

**Where in the World is the Internet?
Locating Political Power in Internet Infrastructure**

by

Ashwin Jacob Mathew

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Information

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor John Chuang, Co-chair
Professor Coye Cheshire, Co-chair
Professor Paul Duguid
Professor Peter Evans

Fall 2014

**Where in the World is the Internet?
Locating Political Power in Internet Infrastructure**

Copyright 2014
by
Ashwin Jacob Mathew

This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike
4.0 International License.¹

¹The license text is available at <http://creativecommons.org/licenses/by-nc-sa/4.0/>.

Abstract

Where in the World is the Internet?
Locating Political Power in Internet Infrastructure

by

Ashwin Jacob Mathew

Doctor of Philosophy in Information

University of California, Berkeley

Professor John Chuang, Co-chair

Professor Coye Cheshire, Co-chair

With the rise of global telecommunications networks, and especially with the worldwide spread of the Internet, the world is considered to be becoming an information society: a society in which social relations are patterned by information, transcending time and space through the use of new information and communications technologies. Much of the popular press and academic literature on the information society focuses on the dichotomy between the technologically-enabled virtual space of information, and the physical space of the material world. However, to understand the nature of virtual space, and of the information society, it is critical to understand the politics of the technological infrastructure through which they are constructed. In this dissertation, I study the infrastructure of the Internet to examine the processes through which the dichotomy between virtual space and physical space is produced. It is through Internet infrastructure that the seemingly placeless appearance of virtual space is produced over disparate physical telecommunications infrastructures which are fixed in space, such as copper telephone wires, satellites, and optical fiber cables.

My examination of Internet infrastructure focuses on the system of interconnections amongst the networks which make up the Internet, which is called the inter-domain routing system. For all that the Internet is spoken of as a singular entity, it is in fact a complex distributed system of over 47,000 interconnected networks spanning the world. It is these interconnections - in the shape of the inter-domain routing system - which allow the Internet to appear to be a single entity, and provide the means through which the apparent placelessness of virtual space is produced. I approach the problem of understanding the production of virtual space by examining the mechanisms involved in the maintenance of order within the Internet's inter-domain routing system.

To examine the mechanisms involved in maintaining order in the inter-domain routing system, I study the technology and practices involved in the interconnection of networks. The technology which enables network interconnection is the Border Gateway Protocol (BGP), which allows the establishment of network interconnections without any need for centralized

oversight. In the absence of centralized oversight, I found that the practices involved in operating BGP rely on coordination and collaboration amongst the technical personnel responsible for managing the interconnection of networks. Coordination and collaboration are enabled amongst the Internet's technical personnel through social relationships of trust, running across corporate and state boundaries. Even though the inter-domain routing system operates without centralized oversight, it does rely on centralized institutional structures for specialized functions, such as standards-setting activity, and the allocation of unique numbering resources required to identify networks in the inter-domain routing system, and to identify computers within networks. The personnel involved in managing network interconnections, and in administering centralized institutions do so - to varying degrees - with an ideal of serving the common good, acting "for the good of the Internet".

I argue that order is maintained within the inter-domain routing system through a distributed system of trust relationships, which are anchored by centralized institutional structures. As an arrangement of mechanisms for maintaining order, I consider this to be a governance arrangement, which I term "distributed governance".

Distributed governance is an unusual, and possibly unique, model of governance. It has three distinguishing features which mark it off from hierarchical and market-based models of governance. First, in its reliance on a distributed system of trust relationships. These are produced and reproduced in the practice of interconnecting networks, and through professional communities of the technical personnel responsible for managing network interconnections. Second, in its centralized institutional structures, which are uniquely organized amongst global governance institutions. None of these centralized institutional structures are formed by international treaty, and all of them are strongly committed to openness and participation. Third, in its operation over the particular technological form of BGP which emphasizes coordination and collaboration. To change the technology of inter-domain routing would be to change the range of governance possibilities for inter-domain routing, modifying the nature of distributed governance itself.

These distinguishing features are sites of contestation. Although technical personnel do owe allegiance to their professional communities, and to one another through trust relationships, they are also employees of corporations which invest in Internet infrastructure, and citizens of nation states which regulate (and sometimes invest in) Internet infrastructure in their territories. Distributed governance is accordingly complicated by market relationships, the interests of nation states, and international relations amongst nation states, just as markets, nation states and international relations are complicated by distributed governance.

To make sense of distributed governance as a global system, I study its instantiation in professional communities of the Internet's technical personnel, centralized institutional arrangements, and state and market interests across two different regions: North America, which is relatively central to the global Internet, and South Asia, which is relatively peripheral. This provides the opportunity to perform a comparison between these two cases, to understand at once how distributed governance varies under different conditions, and how different articulations of distributed governance are linked into a single global system of governance.

The range of social possibilities within a society are shaped by the model of governance which provides it with order. To understand the nature of the information society, it is essential to understand the mechanisms of distributed governance. Indeed, I argue that the social values of “freedom” and “democracy” which are often ascribed to the Internet are only made possible through distributed governance. Equally, these social values cannot be maintained through technology alone. They must be served by an appropriate combination of technology and social arrangements of technical personnel who act with these social values in mind, “for the good of the Internet”.

Contents

Acknowledgments	iii
List of Figures	v
List of Tables	vi
1 Introduction	1
1.1 Where in the World is the Internet?	1
1.2 Locating Political Power in Internet Infrastructure	2
1.3 Mediating Between Physical and Virtual Space	4
1.4 Understanding Distributed Governance	5
1.5 Outline	7
2 Concepts and Methods	9
2.1 Concepts	9
2.2 Methods	21
3 A Social History of the Internet	27
3.1 Sociability in the Early Internet	28
3.2 Building the ARPANET	30
3.3 The TCP/IP Protocol Suite	33
3.4 Building the NSFNET	36
3.5 Networking the Networks	39
3.6 From NSFNET to Internet	42
3.7 Institutionalizing the Internet	46
4 Understanding the Inter-Domain Routing System	47
4.1 The Border Gateway Protocol	48
4.2 Risk and Uncertainty in the Inter-Domain Routing System	52
4.3 A Collective Action Problem	68
5 Distributed Governance I: Institutional Anchors	70
5.1 Encountering the Internet	70

5.2	Developing Standards	72
5.3	Managing Resources	82
5.4	Numbering Networks	94
5.5	Opening Resource Management	109
5.6	Open Institutions, Open Governance	110
6	Distributed Governance II: Ordering Topology	113
6.1	Seeing Dots	113
6.2	Growing the Internet	115
6.3	Going to NANOG Meetings	127
6.4	Becoming a Network Administrator	144
6.5	Producing Interconnection	157
6.6	Operating the Internet	178
7	Distributed Governance III: Ordering Geography	181
7.1	Disembedding the Internet	181
7.2	Extending the Internet	183
7.3	State, Market and Internet	207
7.4	Re-embedding the Internet	216
8	Conclusion: For the Good of the Internet	219
8.1	Trust in Technology	221
8.2	The Production of Virtual Space	230
8.3	Distributed Governance for Internet Infrastructure	237
8.4	Infrastructural Power in the Information Society	240
	Afterword	243
	Glossary	246
	Bibliography	252

Acknowledgments

Since this is a dissertation about infrastructure, it is only fitting that I acknowledge the infrastructure which supported me as I worked through this dissertation. Unlike the infrastructure of the Internet that I write about, which only becomes visible upon failure, the infrastructure of my dissertation was always visible, peopled by a wide variety of individuals, many of whom may never know how much their support meant to me. The infrastructure of my dissertation was a constant and reassuring presence which never failed me as I negotiated the work of research and writing, through various twists and turns, dark alleys and dead ends, and occasional flashes of inspiration.

All of those I spoke with for my research were unfailingly gracious. They welcomed me into their offices and homes, and the social gatherings of their professional communities. I cannot thank them enough for their willingness to take the time to explain their work to me, tolerating my naiveté as I sought to understand the complex world of network operations on the Internet.

Coye Cheshire and John Chuang, my two primary advisers, have been part of my academic life ever since my very first semester at Berkeley. Through classes that I took, and taught with them, and extensive conversations, they taught me how to approach and think through research. They waited patiently when they felt I needed the time to work ideas out, helped me back on track when I lost direction, and gently pushed me along through the dissertation process. I could not have worked my way through this dissertation without their guidance and support.

I learned how to think critically about information from Paul Duguid. Paul was generous, almost to a fault, with his time and advice, as I struggled to make sense of my research. My time at Berkeley was so much the richer for the conversations that I had with him. Paul guided me towards viewing my work not just in interdisciplinary terms, but as research which engages with the concept of information. If I have come to have a better sense of what it means to study information, it is because of Paul.

Peter Evans taught me how to approach problems of globalization. He was constantly enthusiastic and encouraging as I set about the task of researching and writing this dissertation. Every conversation I had with Peter was seminar in miniature, as he helped me focus my work, and bring greater clarity to my account of Internet infrastructure as a global system.

Courses with Gillian Hart and Jean Lave bracketed my theoretical development at Berkeley. Gill welcomed me into her graduate seminar in my first year on campus, even though I was woefully unprepared for it. The deep dive I took into social theory with her provided an invaluable foundation for my later academic work. In particular, she introduced me to the literature on space, which was formative to my thinking about the Internet. Jean Lave taught me how to think about social theory through a class which she co-taught with Paul Duguid. It was through my exchanges with Jean that I gained a clearer perspective on how to approach social theory as an object of study; and on how such an approach can inform the process of research and writing.

The School of Information offered a wonderful and supportive intellectual home. Meg St. John kept the bureaucracy at bay, and was always available for advice throughout my time at Berkeley. Lety Sanchez, as everyone knows, cheerfully keeps the school running. I am thankful to both of them, and to the staff at the school, for providing the administrative support which allowed me to stay focused on research.

I was privileged to be part of a group of doctoral students who became friends and intellectual partners as we worked our way through our research, and collaboratively made sense of what it means to study information. I am thankful to Megan Finn, Elisa Oreglia, Rajesh Veeraraghavan, Janaki Srinivasan, Liz Goodman, Dan Perkel, Daniella Rosner, Christo Sims and Judd Antin, for conversations which shaped my thinking, and for providing mutual support as we navigated our way through our dissertations. I am grateful to Jen King, Sarah Van Wart, Elaine Sedenberg, Robert On, Laura Devendorf, Sebastian Benthall, Bob Bell, Andy Brooks, and everyone else down in the Ph.D. offices for offering ongoing encouragement and conversation in the otherwise solitary and difficult task of writing a dissertation. Of the many students in the broader Berkeley community who helped me think about my research, I must especially thank Greta Marchesi. We met early in our time at Berkeley, and she has been a dear friend and academic partner ever since.

Even though they didn't quite understand what I was doing, my family took pride in my dissertation. I am thankful to them for being there for me throughout my years at Berkeley. Countless friends provided support throughout the doctoral process. My thanks go to Gita, Somu, Azadi, Aman, Albie, Christine, Jeju, Nadia, Kathleen, Shawna, Tom, Nick, Star, Elza, Wolf, Laszlo, Tamar, Ry, Deepak, Aruna, Sunaina, Steve, Jonatthan, Sonja, Robert, Melissa, and so many others. To the kindly folks at Nabolom: thank you for offering a friendly environment with coffee and croissants. When I was stuck in my writing, a morning at Nabolom never failed to get me going again. There are many more people who made this dissertation possible, and I'd like to thank them all, even if I cannot mention them all here.

I am grateful to the TRUST center for funding a year of my research through the National Science Foundation award CCF-0424422, and for providing the opportunity to present my work at the annual TRUST conference.

The research and writing for this dissertation was conducted almost entirely using free and libre open source software. I am eternally grateful to the many volunteers who produce this software, offering an alternative to proprietary operating systems and tools. My thanks also go to Creative Commons, which provides licensing schemes that encourage the sharing of creative work. It is through their efforts that I am able to make this dissertation available under an open license.

If there is a keystone in the infrastructure of my dissertation, it is my wife, Anasuya. When I first waveringly contemplated a return to academia, she was the first to encourage me. Once I got admission to Berkeley, she uncomplainingly pulled up roots from Bangalore, and moved halfway around the world with me. When I doubted myself, she was there to reassure me. When I chanced upon a moment of insight, she celebrated it with me. She read every single word I wrote, debated my ideas with me, and helped me become a better thinker and writer. More than for anyone else, this dissertation is for her.

List of Figures

4.1	Autonomous system relationships	51
4.2	Size of default-free zone over time	60
5.1	ICANN organization chart	86
5.2	The view from audience at an ICANN meeting	88
5.3	Microphone in aisles at ARIN	104
6.1	The view from the audience at a NANOG meeting	129
6.2	Internet connectivity schematic for NANOG57	130
6.3	Total and new attendee counts at NANOG	134
6.4	The proportion of new attendees at NANOG	134
6.5	The core of NANOG attendees	135
6.6	Churn in the core of NANOG attendees	135
7.1	At the SANOG18 routing workshop.	197

List of Tables

7.1	SANOG attendee records by country	190
-----	---	-----

Chapter 1

Introduction

1.1 Where in the World is the Internet?

With the rise of global telecommunications networks, and of the Internet in particular, it is said that we are in the midst of a transformation from industrial society to “information society”, a transformation that is considered to be as profound as that from agrarian society to industrial society.¹ The information society is characterized on the one hand by emancipatory claims that freer flows of information will support democracies, by enabling citizens to hold their governments accountable, and bring about prosperity through new opportunities offered by knowledge work, and integration into market economies through new information services. On the other hand, fears of surveillance and control abound in the information society, especially following Edward Snowden’s revelations about the activities of the NSA and other state intelligence services. However, the exact nature and implications of the information society remain open to question; some even doubt that the information society is a meaningful construct (Garnham 2000; Drori 2007). As the key infrastructure through which the information society is realized, the Internet presents an important site through which to understand the nature of the information society.

Powerful infrastructural technologies - water, gas, electricity, transportation, communications - are integral to modern life, yet are really only noticed when they fail (Star 1999; Edwards 2003). The Internet is the most recent of these, and also the most distinctive: traffic congestion can be experienced on roads, a broken water main is visible and locatable, but it is near impossible to meaningfully conceptualize flows of data on the Internet. The virtual space of the Internet appears to be everywhere - connecting people across the world - and nowhere - somehow impossible to locate.

¹Emblematic of this perspective is the World Summit on the Information Society (WSIS) organized by the United Nations, held in 2003 and 2005. See <http://www.itu.int/wsis/index.html>, last retrieved Aug 2, 2014. See Mattelart (2003) for an overview of theories of the information society, which has also been characterized as “post-industrial society” (Bell 1973) and “post-modern society” (Lyotard 1985; Harvey 1990).

When political power is located in the context of the Internet, the focus is often on phenomena unfolding at the “ends” of the Internet, involving Internet users and Internet content and application providers. Important debates emerge from these studies, engaging with issues including free speech, intellectual property, security and privacy. However, these debates are themselves made possible only through Internet infrastructure, where a disjuncture is produced between the physical space of infrastructure and the virtual space of new media. Many substantive policy conversations are beginning to pay attention to Internet infrastructure, but theories of the information society have largely failed to do so. At best, these theories comment on the manner in which social relations are stretched across space through the use of new information and communication technologies. At worst, they conflate the properties of these technologies with properties of social relations, to argue that technologically-enabled social relations transcend time and space.

If we are to understand the nature of the information society, it is essential to analyze the processes through which the seemingly placeless nature of the virtual space of the Internet is produced. In this dissertation, I approach this problem by examining the related technological and social arrangements of Internet infrastructure, which are integral to the production of virtual space over physical space. In particular, I study the practices of the technical personnel involved in managing the system of interconnections amongst networks - the “inter-domain routing system” - which together constitute the global Internet. This is a critical site for study: in the absence of these interconnections there would be no Internet, and it is only through the interconnection of networks in different geographies that the seemingly placeless experience of virtual space is produced. This dissertation is the first effort aimed at studying the practices involved in operating the inter-domain routing system.

I argue that the infrastructure of the Internet - of which the inter-domain routing system is a critical component - is stabilized and ordered through a system of social relations of trust amongst the Internet’s technical personnel, which cuts across corporate and state boundaries. This system of social relations - which I call “distributed governance” - is essential to the governance of Internet infrastructure, as a distributed system of social relations to manage the distributed technological system of Internet infrastructure. To understand the nature of the information society, the nature of distributed governance must be studied, as the system of social relations which stabilizes and orders the Internet, which is the key infrastructure of the information society.

1.2 Locating Political Power in Internet Infrastructure

The nature of distributed governance in the inter-domain routing system may be clarified with a brief vignette. An unusual event unfolded on the Internet for about two hours on February 24th, 2008, during which all traffic directed to YouTube was sent to a server in Pakistan. The Government of Pakistan had issued an order to censor a particular video on

YouTube, and the Pakistan Internet Exchange (PIE) - a central point of connectivity for many networks in Pakistan - chose to implement this order by issuing a claim to its connected networks that it knew of a better route to YouTube than that which they had. Data traffic along this route was then to be redirected to a web page indicating that YouTube was blocked. Unfortunately, this claim was also transmitted to PIE's Hong Kong-based upstream network provider, PCCW, and from there on to the rest of the Internet. The result was that all traffic directed to YouTube from around the global Internet was redirected to PIE. In response, PCCW disconnected PIE while the rest of the Internet recovered from this fault. Effectively, the bulk of a country's Internet connectivity was disabled through private administrative action.

This event revealed vulnerabilities one of the key technologies of the Internet, the Border Gateway Protocol (BGP), which enables the interconnection of networks to form the global Internet. PIE issued their spurious claim about the location of YouTube using BGP, and it was through BGP that this claim spread to the rest of the Internet. It seems unlikely that an entity in Pakistan would be able to hijack traffic directed to one of the Internet's most significant and visible websites, YouTube, operated by one of the most important corporations on the Internet, Google; and yet it happened. Such failures are, in fact, well recognized and not uncommon, occurring on an ongoing basis at varying scales. What is remarkable about the Internet is that it continues to provide the appearance of stability to ordinary users in spite of such vulnerabilities.

The appearance of stability of the Internet is made possible by the distributed relationships of coordination and collaboration amongst the technical personnel who administer Internet infrastructure, which were also revealed by this event. Google personnel were able to contact the relevant personnel at PCCW and have them take the necessary action to disconnect PIE, even in the absence of a contractual relationship with Google requiring PCCW personnel to do so. A stable and widespread system of social relations for coordination and collaboration amongst the Internet's technical personnel - which cuts across corporate and state boundaries - is essential to maintaining the stability of the Internet.

From the perspective of PCCW personnel, they were disconnecting a single errant entity, PIE. Yet this disconnection effectively implied the disconnection of a country from the Internet. To be clear, this was not an anti-censorship action against the Pakistan government, but rather a response to the manner in which PIE chose to implement the Pakistan government's censorship order, and the mistake of technological configuration that PIE made, which propagated their censorship to the rest of the Internet. This event illustrates how the system of relations amongst the Internet's technical personnel may be conceived of as a political actor in its own right, which may under certain conditions counterbalance the power of sovereign governments. As I will show, this system of relations can also act autonomously in relation to the economic interests of the corporations which employ these technical personnel. Coordination and collaboration across corporate boundaries can require a degree of openness and transparency which may run counter to the economic interests of corporations.

In this dissertation, I explore the politics of the Internet, focusing not on individual services such as YouTube, but rather on the infrastructure of the Internet which allows these

services the appearance of being simultaneously everywhere and nowhere. I focus on the related social and technological mechanisms through which the appearance of separation between virtual and physical is achieved through an examination of the Internet's inter-domain routing system: the system of networks interconnected with BGP to form the Internet.

I will show how the distributed system of relationships amongst technical personnel stabilizes and orders the distributed system of interconnections amongst networks in the inter-domain routing system. I argue that these arrangements constitute a distinctive system of governance - "distributed governance" - that is based upon trust relationships amongst technical personnel, enacted in the everyday practices of network administration, and produced and reproduced across time and space through the Internet's technical communities.

1.3 Mediating Between Physical and Virtual Space

The infrastructure which I reference here is not the physical telecommunications layer of optical fiber cables and satellites, but rather the "logical layer" of the Internet, which mediates between physical and virtual space. The Internet is made possible through the Internet Protocol (IP), and related technologies such as BGP, which construct the logical layer that interconnects various physical telecommunications infrastructures - including satellites, submarine cables, WiFi and phone lines - to present a single global system over which services in virtual space - such as YouTube - may be constructed.

It is essential to understand the means through which shifting conceptions of space are produced across different layers of Internet infrastructure, following the processes through which the logical layer translates between physical and virtual space. If we are to form a theory of power in the information society, it must necessarily engage with the manner in which power is structured within and across the different layers of Internet infrastructure.

States are central to this analysis at the physical layer. Within national territory, physical infrastructure has always been strongly regulated and controlled by the state, from managing allocation of wireless spectrum to granting rights to lay copper or optical fiber cables. State authority also remains paramount for physical infrastructure deployed across international spaces to provide global connectivity. The United Nations Convention on the Law of the Sea specifies rights and obligations for states with regards to submarine telecommunications cables laid in their territorial waters, and on the high seas.² Similarly, the International Telecommunications Union allocates orbital slots for telecommunications satellites only to states. In order to obtain rights to deploy physical infrastructure, corporations must engage with nation states, and international bodies in which states maintain control.

The logical layer of network infrastructure, which emerges with the Internet, inverts these power relations. As the Pakistan/YouTube incident illustrates, the Internet's technical personnel can act autonomously in relation to nation states - and even in relation to the corporations employing them - in the interests of maintaining the stability of Internet

²See http://www.un.org/Depts/los/convention_agreements/texts/unclos/part7.htm, last retrieved Jul 24, 2014.

infrastructure. A similar inversion of power obtains in the centralized institutions of Internet governance - responsible for standards setting activity, and the allocation of resources - in which technical and economic interests dominate the interests of nation states.

While I analytically treat nation states, corporations operating Internet infrastructure, and the Internet's technical personnel as distinct entities, this is not to say that they are distinct in reality. They overlap, interpenetrate, and shape one another in a variety of ways, over the terrain of the technology of the Internet's infrastructure. The distributed governance of Internet infrastructure must, therefore, be understood in its relation to nation states and corporations, as much as in its own distinctive internal logic.

1.4 Understanding Distributed Governance

The general problem that I take on in this dissertation is the question of how the Internet is produced and maintained as a stable, global infrastructure through distributed social relations, rather than through policing and control by centralized authorities, or through purely economic arrangements. I examine this problem through the lens of the BGP, and the system of interconnections amongst networks which it enables: the Internet's inter-domain routing system. As the Pakistan/YouTube incident illustrates, failures of the inter-domain routing system can affect the entire Internet, making this a critical site through which to understand Internet infrastructure.

I approach this problem with the following research questions:

- What are the mechanisms through which distributed social relations are organized to stabilize and order inter-domain routing, as a unified system of distributed governance?
- What is the relationship between technological form and forms of governance?
- How does the logic of distributed governance interact with the interests of corporations and nation states?

To address these questions, I explore three related propositions, which are based on the interactions of nation states, corporations, and the Internet's technical communities, organized in terms of trust, space, and technological form.

First, I argue that *trust* is the essential mechanism that enables coordination and collaboration in distributed social relations amongst the Internet's technical personnel. In fact, as I will show, trust relations amongst the technical community of the early Internet strongly contributed to the design of BGP as a protocol which assumed trust relations in the practice of inter-domain routing. The risks and uncertainties inherent in BGP are carried across space and time through the inter-domain routing system in the technological form of BGP itself. Trust relations for handling these risks and uncertainties are part of the everyday *practice* of inter-domain routing, anchored by centralized governance institutions, and produced and reproduced through regionally organized technical *communities*.

Although trust, practice and community can explain the internal logic of distributed governance, they are insufficient to examine the interaction of distributed governance with the interests of corporations and nation states. The concept of *embeddedness* - of markets for network interconnection, and of centralized governance institutions - provides means through which to understand this interaction.

Second, I argue that the elements of distributed governance - trust, practice, and community - are produced and reproduced through face-to-face meetings of the Internet's technical communities, just as much as these elements are mutually constitutive of the *space* of Internet infrastructure. To stabilize and order a global Internet, the logic of distributed governance must be extended over space. Equally, the space of Internet infrastructure is produced through the activity of distributed governance, which must be understood in terms of the elements of trust, practice and community, as well as in terms of the processes required to extend these elements across space to produce the space of the global Internet.

This space, however, is not produced through distributed governance alone; corporations invest capital in the production of this space, and nation states regulate - and sometimes invest in - Internet infrastructure in their territories. The processes involved in the production of the space of Internet infrastructure provide additional insights into the possibilities and limits of distributed governance, as much as they do into the ways in which nation states and corporations shape - and are shaped by - distributed governance.

Third, I argue that the specific technological form of BGP encodes within itself a particular imaginary of power, which assumes a system of infrastructure which is stabilized and ordered through coordination and collaboration amongst autonomous entities, rather than through policing and regulation by a centralized authority, or through market relations. This is not to say that centralized authorities and market relations do not play a role in the inter-domain routing system. Rather, I argue that the technological form of BGP offers the basis upon which the power of the Internet's technical communities is constructed, while at the same time limiting the power of corporations and nation states. If this technological form were to change, then the possibilities for governance - and the balance of power - within the inter-domain routing system would also change.

I adopt qualitative methods for this research, performing a relational comparison between professional communities of technical personnel responsible for inter-domain routing in North America and South Asia, alongside research into centralized governance institutions which perform standards setting activity and manage resources for Internet infrastructure. The study of centralized governance institutions is required to understand the role they play in anchoring distributed governance. The relational comparison between the North American and South Asian contexts is essential to understand distributed governance as a global system, instantiated in different regions with different local conditions, resulting in different regional outcomes; but still related to one another through technology, practice and social relations to ensure the global stability of the Internet. The choice of these particular contexts for research is dictated by their relative positions: North America is relatively central to the global Internet, while South Asia is relatively peripheral. Accordingly, the choice of these particular contexts offers the opportunity to establish the different concerns of, and

the connections between, central and peripheral positions towards forming a more complete perspective on distributed governance.

Although the history and centralized governance institutions of the Internet have been well studied, little attention has been paid to the practices of technical communities, which I focus on here. The study of these technical communities is essential to forming fuller understanding of Internet governance, especially in relation to their involvement in the operation of Internet infrastructure, and the Internet's inter-domain routing system in particular. This research extends understandings of Internet governance to present a perspective on governance - "distributed governance" - which focuses on the distributed social system required to manage the distributed technological system of Internet infrastructure.

Technology has long been considered to be a force which shapes society, possessed of a logic of its own. Marx (1964) calls this kind of thinking "the technological sublime", the notion that new technologies will lead to social betterment, as they reduce spatial and temporal barriers to social integration. Marx was writing of the railroad and steamships, but the same kind of thinking is apparent in theories of the information society, which add information as a force for change alongside technology. By illustrating the social work that goes into producing Internet infrastructure, I aim to dispel the myth that technology and information - especially in the shape of the Internet - are autonomous forces shaping society, offering inherent emancipatory possibilities for "freedom" and "democracy", or functioning as a perfect means of surveillance and control. It is the complex of socio-technical relations of distributed governance of Internet infrastructure - and its interaction with established political and economic powers - which shape the social possibilities that the Internet provides, making this a critical area of study.

1.5 Outline

This dissertation consists of three parts, each of which I discuss in turn here: background, empirical analysis, and theoretical implications.

In the first part of this dissertation, consisting of Chapters One through Four, I elaborate my research problem, and provide the historical and technical background necessary to understand my argument.

Chapter Two provides an overview of the theoretical perspectives that I adopt in this dissertation, and details my research methods. The two principal areas of theory which I draw from are of trust and space, alongside theories of embeddedness and community. In addition, I present an original analytical framework for research into Internet infrastructure.

Chapter Three offers an account of the early history of the Internet. There are many other histories of the Internet (several of which I draw from for this chapter), but my account is distinctive in its focus on the arguments presented in Chapter One, and in developing the analytical and theoretical perspectives from Chapter Two.

Chapter Four presents an extensive overview of the Internet's inter-domain routing system, which is the primary subject of my research. I discuss the mechanics of the Border

Gateway Protocol in detail in this chapter, along with the various risks and uncertainties involved in the operation of the inter-domain routing system.

In the second part of this dissertation, consisting of Chapters Five through Seven, I present the empirical material which provides the basis for my analysis of distributed governance in Internet infrastructure. Chapters Five and Six detail the internal logic of distributed governance, drawing largely from the context of the USA, while also dealing with the interaction between distributed governance and economic interests. Chapter Seven presents empirical research from South Asia, as a relational comparison with the material presented in Chapters Five and Six, to construct a more complete perspective on the work involved in constructing distributed governance as a global system.

Chapter Five describes the key centralized governance institutions which anchor distributed relations of governance. The institutional anchors discussed in Chapter Five are the Internet Engineering Task Force (which sets technical standards for Internet infrastructure), the Internet Corporation for Assigned Names and Numbers (the global root authority for the management of Internet resources), and the American Registry for Internet Numbers (the authority managing number resources, such as IP addresses, in North America).

Chapter Six focuses on the manner in which distributed relations constitute a system of governance for the Internet, anchored by the centralized institutions discussed in Chapter Five. The role of trust in enabling coordination and collaboration in the practice of network administration is discussed, as is the way in which this distributed system of trust relations interacts with the economic interests of the private entities operating Internet infrastructure. Through this discussion, trust is presented as the key mechanism through which distributed governance operates, and technical communities are presented as the site through which trust relations and technical practices are produced and reproduced.

Chapter Seven discusses the extension of distributed governance to South Asia, and the limits of the logic of distributed governance in its interaction with political and economic interests in India. The problem of space is central to this chapter, both in terms of understanding the extension of trust relations for distributed governance across space, and in terms of understanding the spatial logic of distributed governance in its interaction with the territorial spatial logic of the Indian state.

In the final part of this dissertation, consisting of Chapter Eight, I discuss the theoretical implications of distributed governance. I draw together theories of trust and space to provide a unified theoretical perspective on distributed governance, and its limits. I close with a discussion of the implications of distributed governance for theories of the information society.

Chapter 2

Concepts and Methods

In this chapter, I review some of the basic concepts that I use in this dissertation: trust, embeddedness, community, space, and infrastructure. Trust is the essential social relation over which the distributed governance of the Internet's inter-domain routing system is constructed. Trust relations are produced and reproduced in technical communities, which also play a key role in extending trust relations across space. I use the concept of embeddedness to examine the interaction of distributed governance interacts with states and corporations. I draw from theories of infrastructure to form my conception of the inter-domain routing system as infrastructure.

Following the review of concepts, I comment on the methodological issues I faced in designing and conducting this research, and I detail the research strategy I adopted in consequence. In particular, I am concerned with how to think about a protocol - the Border Gateway Protocol - as a field site for ethnographic research.

Finally, I detail the specifics of the data I collected, including textual sources, field sites and interviews.

2.1 Concepts

2.1.1 Trust

The essence of trust lies in the management of expectations. These may be the expectations of everyday life - such as expecting that water will flow when a tap is opened - making trust co-extensive with social reality. Alternatively, these may be more narrowly scoped expectations in interpersonal interactions - such as trusting a plumber to fix a faulty tap - making trust but one element of social reality. Theories of trust emphasize each of these perspectives in varying degrees.

Trust is considered to be a response to problems of risk and uncertainty. Risk reflects what is at stake (such as the amount of money in an investment), while uncertainty is the probability that a risk taken will not result in positive returns (Cook et al. 2005). Trust is

also considered to be a response to complexity (Luhmann 1979) and scale (Giddens 1991) in the modern world. Luhmann (1979) calls this “system trust”, to distinguish it from the interpersonal trust we may have in one another.

It is remarkably difficult to pin down the nature of trust, even though the problems which trust addresses are well-defined. As Luhmann puts it:

Trust rests on an illusion. In actuality, there is less information available than would be required to give assurance of success. The actor willingly surmounts this deficit of information. (Luhmann 1979:32)

Lewis and Weigert (1985) make a similar point, when they consider rational prediction as an alternative to trust. They point out that it is not possible to gather and process sufficient information to reliably predict every possible outcome of an action. Trust is the social mechanism which supports belief in certain desirable outcomes, in the face of complexity, risk and uncertainty.

In their definition, Lewis and Weigert argue that trust is composed of three dimensions: cognitive, emotional and behavioral. One who would trust (the truster) makes a cognitive assessment of the object of trust (the trustee), to ascertain whether the trustee is trustworthy. If the trustee betrays the trust reposed in them, emotional damage results for both truster and trustee; similarly, when trust is honored, both truster and trustee experience positive affect. Once trust has granted, the truster behaves as though the future ensured by trust will come to pass. This helps expand the basis for the cognitive dimension of trust, as others observe the truster’s behavior, and the trustee’s actions, allowing them to establish the a sense of the trustworthiness of the trustee.

Powell (1990) echoes this last point in his discussion of network forms of organization, arguing that conflict resolution in networks is governed by reciprocity and reputation. Following Lewis and Weigert (1985), this implies widespread trust relations, and the ability to observe - or at least receive reports about - one another’s behavior. Widespread trust relations may also contribute to, and be constructed through, general trust within a society. Yamagishi and Yamagishi (1994) show Americans to have higher general trust (i.e., a greater propensity to trust one another) than Japanese. They use their results to distinguish between trust, and cooperation enabled through assurance structures, which are constructed through networks of stable committed relationships. They argue that “trust is based on the inference of the interaction partner’s personal traits and intentions, whereas assurance is based on the knowledge of the incentive structure surrounding the relationship.” (Yamagishi and Yamagishi 1994:132).

Assurance structures reduce risk and uncertainty, supporting cooperative behavior without any need for trust. The contrast between assurance and trust leads Gellner (1988) to articulate the trust paradox:

... it is precisely anarchy which engenders trust or, if you want to use another name, which engenders social cohesion. It is effective government which destroys trust (Gellner 1988:143).

Gellner argues that it is in the interests of government (an assurance structure) to maintain atomized social relations, ensuring its continuity by requiring citizens to depend on it for the maintenance of social order. Assurance structures may, therefore, be conceived of as a stable institutional arrangements which anchor cooperative behavior, whether centralized (as in government), or distributed in social norms and networks of committed relationships. When considered from the perspective of social order - which is Gellner's concern - a strong assurance structure may result in weak trust relations; similarly, a weak assurance structure creates the conditions for the emergence of strong trust relations.

The concept of networks of committed relationships as assurance structures leads to a consideration of the relationship between trust and choice. Yamagishi and Yamagishi (1994) argue that trust is required in order to make the choice of switching from a committed relationship to a new relationship. Luhmann (1988) offers a more refined perspective on this point, arguing that choice is a prerequisite for trust. If one party has no option but to cooperate with another party their relationship may be best characterized as confidence in Luhmann's account. In contrast, where one party can choose whether or not to cooperate with another party, then trust is required to enable cooperation.

Just as socio-cultural contexts vary, so does the nature of trust vary. Trust, therefore, is learned behavior. From a broad perspective, behaviors associated with trust - how to trust - are learned through socialization (Luhmann 1979:27-28). From a narrower perspective, trust relations are formed through repeated interactions. An individual's trust in another is formed and grows as they repeatedly engage in cooperative behaviors with successful outcomes. In this perspective, a one-shot interaction cannot be characterized in terms of trust, but must rather be thought of simply as a gamble (Hardin 2002:14-17). Repeated cooperative behaviors may lead to thick relationships with commitments on both sides; but thick relationships are not definitive of trust, as Hardin (2002) points out:

The correct way to see thick relationships, however, is as one possible source of knowledge for the truster about the trustworthiness of another and one possible source of incentives to the trusted to be trustworthy (Hardin 2002:22).

Luhmann (1979) argues that assurance structures are symptomatic of modern societies, so much so that the very nature of trust has changed. Luhmann is concerned with the problem of the reduction of complexity required in the course of everyday life in the modern world. Such complexity reduction, he argues, is accomplished through trust in the systems which enable our lives - banking, electrical supply, roadways, and so on - which he terms "system trust". In his account, system trust supplanted interpersonal trust as the dominant model of trust in the progression from simpler pre-modern settings to the complexity of modernity. Giddens (1991) makes a similar point, although his concern is with the mechanisms of globalization, which allow social relations to be extended over space and time. He argues that this is accomplished through trust in expert systems (such as the expertise of engineers or pilots) and trust in symbolic tokens (such as money). Giddens characterizes such trust in abstract capacities in terms of faith or confidence (Giddens 1991:27).

In contrast, Hardin (2002) considers trust to always be an interpersonal relation, arguing that the relation which obtains between individuals and institutions - as between citizens and their governments - cannot be characterized as a trust relation. He is similarly skeptical of accounts of general trust. To make his point, Hardin distinguishes between trust and trustworthiness, arguing that many perspectives on trust are better conceived of as trustworthiness. He treats trust as a three part relation, in which the *truster* trusts the *trustee* to perform a specific *action*. In his conception, trust is not a general relation between truster and trustee, but must be understood in terms of the object of trust, an action. Through this three part relation, the trustee encapsulates the interests of the truster in their own interests, leading to Hardin's definition of trust as "encapsulated interest". This provides the basis for a wide-ranging critique of theories of trust. For instance, in Hardin's account, general trust may be understood as a two part relation (a trustor trusts a trustee for any arbitrary action), or even a one part relation (a trustor trusts anonymous others). In the absence of encapsulated interest, Hardin argues, there cannot be trust.

The problem with Hardin's theorization of trust, as Farrell (2009) points out, is that it fails to take into account the broader social structures which may condition the formation of trust relations. I follow Farrell's position in this dissertation, explaining cooperative behavior in terms of a mix of assurance structures and trust relations, never purely one or the other. My concern is to ascertain the relative salience of, and interdependence between, assurance structures and trust relations under different conditions.

2.1.2 Embeddedness

The concept of embeddedness is used to describe the relationship between a system which operates on a seemingly immanent logic (such as a "free" market), and an associated social system. Embeddedness may also be used to discuss the relationship between two social systems which operate on different underlying logics.

Polanyi (2001) originated the concept of embeddedness with his discussion of the embeddedness of markets in society. His concern was with the notion that markets are subject to natural laws of their own, to which society must be subject. He regarded this perspective as promoting the disembeddedness of markets from society, which he argued would destroy society:

Our thesis is that the idea of a self-adjusting market implied a stark utopia. Such an institution could not exist for any length of time without annihilating the human and natural substance of society; it would have physically destroyed man and transformed his surroundings into a wilderness (Polanyi 2001:3).

In Polanyi's analysis, the disembedding of markets from society occurs through the commodification of elements of industrial organization which are not produced for sale. Specifically, these "fictitious commodities" are land, labor and money:

Labor is only another name for human activity which goes with life itself, which in turn is not produced for sale but for entirely different reasons, nor can that activity be detached from the rest of life, be stored or mobilized; land is only another name for nature, which is not produced by man; actual money, finally, is merely a token of purchasing power which, as a rule, is not produced at all, but comes into being through the mechanism of banking or state finance. None of them is produced for sale (Polanyi 2001:75-76).

The attempt to disembed markets from society, says Polanyi, is doomed to failure; since markets are produced through society, they cannot be disembedded from society. However, the very attempt itself causes immense social dislocations, which call forth a counter-movement to re-embed markets in society. The goal of policy, therefore, should be to regulate markets to serve social needs to find the place of markets within broader social relations, rather than to imagine society as being organized entirely through market relations.

Jessop (2007) extends Polanyi's analysis to examine knowledge as a fictitious commodity in what he terms "knowledge-based economies", which I refer to as the information society. He points out that knowledge is collectively produced, and is not inherently scarce; therefore we must examine the processes through which knowledge is made to be valuable. He identifies three such processes, all of which involve the disembedding of knowledge from social relations. First, knowledge is detached from manual labor and material products to construct expert systems and services. Second, knowledge is disembedded from social contexts to be re-imagined through market logic as being profitable or unprofitable. Third, knowledge moves away from circulation through reciprocity and redistribution, and instead is allocated through market systems. States are integral to these processes, as they construct the legal regimes through which knowledge is commoditized. States are also sites of opportunity for counter-movements, providing the means through which knowledge may be re-embedded in society, whether through legal regimes which support intellectual commons, or through renewed investments in public research institutions.

It is worth remembering, as Polanyi (1966) points out, that there is a distinction which must be drawn between explicit knowledge which may potentially be disembedded in the senses that Jessop outlines, and tacit knowledge which cannot be disembedded from social relations. All knowledge is a mix of explicit and tacit knowledge: "we know more than we can tell" (Polanyi 1966:4). To disembed knowledge, to make it explicit and subject to commoditization, is to lose the tacit dimension of knowledge.

Granovetter (1985) offers a perspective on embeddedness which integrates the concept of trust to critique the notion that markets operate logics which are independent of social relations. He points out that conceptions of markets tend to be either undersocialized or oversocialized, and offers trust as a means to understand the mechanisms through which markets are embedded in social relations. He describes an undersocialized perspective as assuming individuals engaging in "rational, self-interested behavior minimally affected by social relations" (Granovetter 1985:481). This is the idealized economic perspective on market exchange, and even on the institutions which evolve to regulate and police market ex-

change, effectively substituting for social relations. An oversocialized perspective construes individuals as “obedient to the dictates of consensually developed systems of norms and values, internalized through socialization, so that obedience is not perceived as a burden” (Granovetter 1985:483). Granovetter characterizes assumptions about the top-down flow of authority in hierarchies, and conceptions of a “generalized morality” in market exchange as examples of oversocialized perspectives.

Both undersocialized and oversocialized perspectives, Granovetter points out, assume atomized individuals, subject to uniform and universal systems of rules, with little attention to social relations. He argues that attention to the role of trust relations in economic exchange offers a path between “the oversocialized approach of generalized morality and the undersocialized one of impersonal institutional arrangements by following and analyzing concrete patterns of social relations” (Granovetter 1985:493).

Evans (1995) develops the concept of embedded autonomy to analyze the relationship between the administrative structures of the state and the society it serves. Administrative structures must, on the one hand, be autonomous to serve the needs of society at large. On the other hand, they must be embedded in state-society relations to allow for involvement from society in the formation of state policy. Evans argues that a state can be developmental only when both conditions of embeddedness and autonomy are satisfied.

These conceptions of embeddedness may broadly be divided into two perspectives. On the one hand, Polanyi (2001) and Jessop (2007) treat embeddedness as a commodity relation, analyzing the embeddedness of markets in terms of the (fictitious) nature of the commodities being traded. On the other hand, Granovetter (1985) and Evans (1995) treat embeddedness as a social relation, thinking of markets and administrative structure in terms of the social relations through which they are constructed, internally and in relation to other social groups.

Giddens (1991) is concerned with the mechanisms that enable the disembeddedness of social relations from space and time in modern societies. He argues that this is made possible through trust (which is closer to confidence, or faith, in his account) in abstract systems, which are either expert systems (of doctors, architects, and so on) or symbolic tokens (such as money). The context for the production of trust may remain abstract and faceless - as in the everyday use of money - or it may be re-embedded through facework - as in encounters with professionals who are part of an expert system.

I make use of these various senses of embeddedness to discuss the accountability of governance institutions, and the embeddedness of systems of technology and information (of Internet infrastructure) in related social systems.

2.1.3 Community

Community is a remarkably difficult concept to define; there are many different concepts of community, covering a range of theoretical perspectives. From the “imagined community” which is not based on everyday interaction, but rather on a recognition of shared symbols through the rise of nationwide media (Anderson 2006), to “epistemic community” which is constituted of networks of experts, encoding certain kinds of knowledge (Haas 1992), to

“virtual” or “online” communities on the Internet (Rheingold 1995; Donath 1999; Wilson and Peterson 2002), and many more.

In this dissertation, I narrow down on two specific theoretical perspectives on community which equip me with the analytical purchase to engage with my research questions. The first is of communities of practice (Lave and Wenger 1991), and the second is of community as symbolically constructed (Cohen 1985). These theoretical perspectives provide the means to understand the relationship between the practice of network administration and the technical communities I studied, and to understand the shared symbolic resources involved in constructing and maintaining these communities across time and space.

Lave and Wenger (1991) argue that understandings of learning cannot be limited to classroom settings, and must be extended to learning as it is situated in broader social relations. They use the term “community of practice” to highlight the ways in which social practice and community produce and reproduce one another. They analyze the process of “legitimate peripheral participation” through which practices are learned in a variety of communities; and in which learning is an integral component of practice. The legitimacy of participation defines “ways of belonging” (Lave and Wenger 1991:35), which provokes attention to the manner in which newcomers are socialized into a community of practice, to processes of becoming oldtimers, and the relations between newcomers and oldtimers (Lave and Wenger 1991:56-57). Peripherality indicates accepted social positions from which newcomers may enter a community of practice, and move towards fuller participation in the community as they acquire greater skills in their practice. This is not to say that there is a single center to which peripheral participants aspire, but rather that there are more central positions defined by social relations within a community of practice, which may be arrived at through fuller participation (Lave and Wenger 1991:36-37).

Cohen (1985) characterizes community in terms of symbolic construction, examining the processes through which the cohesiveness and boundaries of community are constructed through shared symbolic resources. Although symbols may be shared, meaning need not be; individuals in a community ascribe meaning to symbols filtered through their own social experience. For instance, rituals represent symbolic resources which are experienced differently by different participants. Shared symbols provide the means through which the distinction between insiders and outsiders can be constructed in a community; as much as a community is defined through internal cohesion, it is also defined in relation to other social formations outside it.

Through his analysis of the symbolic construction of community, Cohen is concerned with deconstructing several myths about community. First, that community is a simpler face-to-face form of society, which has been succeeded by more complex modern societies; community has existed in various forms across history. Second, that community is inherently egalitarian; all communities are fraught with power relations. Finally, that communities will lose their distinctive character once they are linked up with the infrastructures of modernity, and subject to mass media, consumerism, and so on; rather, ideas are re-appropriated in transiting a community boundary.

When I use the term community, I mean it both in terms of practice and symbolic

construction, but also structurally, as a nexus for the production and reproduction of trust relations. As I show in the following chapters, the practice that I examine depends upon trust relations, and the community that forms in relation to this practice takes on trust as one of its defining characteristics.

2.1.4 Space

Space is often conceived of merely as a container within which human activity unfolds over time. As Soja (1989) puts it in his call for renewed attention to space in social theory:

Space tends to be treated as fixed, dead, undialectical; time as richness, life, dialectic, the revealing context for social theorization (Soja 1989:11).

Rather than treating space as static, Lefebvre (1991) in his landmark work seeks to understand the production of space:

Though [space is] a *product* to be used, to be consumed, it is also a *means of production*; networks of exchange and flows of raw materials and energy fashion space and are determined by it. Thus this means of production, produced as such, cannot be separated from the productive forces, including technology and knowledge, or from the social division of labor which shapes it, or from the state and the superstructures of society (Lefebvre 1991:85).

Lefebvre's conception of the production of space is expansive, linking the material infrastructures and flows produced through human activity to the state, the economy, and the development of technology. His analysis is based upon a theoretical triad of *spatial practice* through which society constructs its space, *representations of space* through which experts (planners, architects, engineers and so on) conceive of space, and *representational space* which encompasses the symbolic means employed in inhabiting space (Lefebvre 1991:38-39).

Through this analysis, he arrives at a perspective on the link between social relations and space, insisting that this link itself must be an object of study; the appearance of particular qualities of space or social relations are produced through the nature of the link between them:

Social relations, which are concrete abstractions, have no real existence save in and through space. *Their underpinning is spatial*. In each particular case, the connection between this underpinning and the relations it supports calls for analysis. Such an analysis must imply and explain a genesis and constitute a critique of those institutions, substitutions, transpositions, metaphorizations, anaphorizations, and so forth, that have transformed the space under consideration (Lefebvre 1991:404).

Harvey (2007) argues that capital can overcome space only by fixing itself in space. Spatial fixes of capital in infrastructure are required to enable flows of transport, energy,

information, and so on, to produce the effect of “time-space compression” (Harvey 1990) in social relations. Harvey also considers space as the means through which capital can resolve its problems. For instance, issues with labor can be spatially fixed by moving manufacturing to a different location.

Following Harvey, it is important to understand the division of labor in spatial terms. Massey (1994c) argues that the object of concern here is “the stretching out over space of the relations of economic ownership and possession”(Massey 1994c:87). She goes on to point out that this incorporates notions of authority - as between a head office and branch office - and specialization - as between the work of design performed in one location, and the work of production in another - as well as relations of market exchange. Her primary concern is with the spatial organization of class relations, which, as she states emphatically, “do not, as the saying goes, exist on the head of a pin” (Massey 1994c:87).

Castells (2000) argues that the dynamics involved in the stretching of relations over space has resulted in a qualitative shift in the nature of space itself. He characterizes this shift in terms of a binary opposition between a networked “space of flows” and a geographically bounded and disconnected “space of places”. The “space of flows” is “the material organization of time-sharing social practices that work through flows” (Castells 2000:442), while the “space of places” is “a locale whose form, function, and meaning are self-contained within the boundaries of physical contiguity” (Castells 2000:453). In the “network society” of the “space of flows”, space-time exists in a binary state of either complete collapse (for nodes within the same network), or an insurmountable infinity (for nodes outside a network):

The topology defined by networks determines that the distance (or intensity and frequency of interaction) between two points (or social positions) is shorter (or more frequent, or more intense) if both points are nodes in a network than if they do not belong to the same network. On the other hand, within a given network, flows have no distance, or the same distance, between nodes. Thus, distance (physical, social, economic, political, cultural) for a given point or position varies between zero (for any node in the same network) and infinite (for any point external to the network). The inclusion/exclusion in networks, and the architecture of relationships between networks, enacted by light-speed-operating information technologies, configure dominant processes and functions in our societies (Castells 2000:501).

The fundamental conflict in the “network society”, Castells argues, is between the differing logics of the “space of places” and the “space of flows”. The logic of the “space of flows” is dominant “because function and power in our societies are organized in the space of flows, the structural domination of its logic essentially alters the meaning and dynamic of places” (Castells 2000:458). However, “there follows a structural schizophrenia between two spatial logics that threatens to break down communication channels in society” (Castells 2000:459).

Whose perspective is more accurate? Castells argues that the dimensions of space-time have collapsed in the “network society”. In contrast, Massey, Harvey and Lefebvre present

positions which suggest that such a collapse of space-time is not possible, and must be understood in terms of its construction, rather than the appearance of its existence. I draw from these latter insights to argue that what is necessary is a clearer understanding of the spatial underpinning which supports the appearance of collapsed space-time in the social relations of Castells' "network society".

Others have critiqued Castells for being insufficiently precise in specifying his key concept of networks, and for a lack of thoroughness in his empirical analysis (Abell and Reyniers 2000; Heiskala 2003). Here, I remain concerned with theoretical responses to Castells which I use to shape my research. In addition to those which I have already reviewed, there are additional perspectives on networks, place, scale and territory which are useful for thinking about the production of space in contemporary societies.

Sheppard (2002) cautions against the image of networks as being constructed through horizontal spatial relationships. Instead, he recommends a focus on relational inequalities in networked spaces, through the analytic of positionality. In his perspective, power must be understood in terms of position as constructed in relation to other positions. This calls for research which is at once comparative and relational in its analysis; cases may be compared to make sense of a larger system only insofar as their relations to one another, and their positionality in the system as a whole are taken in to account.

Massey (1994a) is concerned with rescuing the concept of place, and the notion of community as place-bound, much as was Cohen (1985). Massey argues that place must be conceived of not as an easily bounded static entity, but rather as a nexus of social interactions, reaching into and through a place from and to remote locations, recasting place itself as an ongoing dynamic process. This is not to say that places are doomed to uniformity as they conform to the logic of social interactions reaching over broader scales; rather, the uniqueness of places may be found in the specific mixture of wider and more local social relations which produce place.

The concept of place calls to attention the problem of scale. If place is a local scale, albeit produced through both local and wider social relations, it becomes important to consider scale itself as an analytic for the production of space. Agnew (1994) cautions against the "territorial trap" that comes with treating the state as a unit of analysis, whether as sovereign territorial space, as a container of society, or to construct a domestic/foreign polarity. Treating the state as the primary spatial scale for analysis raises a variety of issues for analysis, not the least of which is that it can result in a blindness to the ways in which state territoriality is produced - and challenged - through interactions between processes unfolding at subnational and supranational scales. Brenner (1999) extends this argument, to conceptualize interactions across multiple scales (including localities, regions, cities and states) in processes of globalization as:

... a dialectical interplay between the endemic drive towards space-time compression under capitalism (the moment of deterritorialization) and the continual production of relatively fixed, provisionally stabilized configurations of territorial

organization on multiple geographical scales (the moment of reterritorialization) (Brenner 1999:43).

Following Brenner, the state cannot be viewed as withering away in the face of processes of globalization, but rather must be understood as an integral element in the unfolding of these processes.

Jessop et al. (2008) attempt to provide an integrated perspective on the use of territory, place, scale and networks in the analysis of space. In their method, each of these analytical elements may be thought of as a structuring principle, which operates over all of these elements considered as fields of operation. For instance, territory as a structuring principle may define territory itself in terms of borders, place as located in a territory, scale as multilevel government, and networks in terms of the inter-state system or state alliances (Jessop et al. 2008:395). I follow their advice in this dissertation, separating territory, place, scale and networks as analytical categories, but still focusing on the ways in which they shape one another.

2.1.5 Infrastructure

The modern world is characterized by infrastructures which are remarkable for the ways in which they are taken for granted, as part of a common “technological unconscious” (Thrift 2004). Edwards (2003) argues that infrastructures offer the means through which technology and society are held to be ontologically separate in modern societies. When an infrastructure breaks down, it seen to be a technical problem; yet little attention is paid to the social formations required to maintain infrastructure, which are themselves part of the societies which use infrastructure. In Edwards’ account, all infrastructure is socio-technical, consisting of technology which is produced by society, which in turn provides the basis for the existence and extension of modern society across time and space. Innis (1951) makes similar propositions in relation to communications media, categorizing communications media as biased either towards time (stable over long time scales, but not easily transported) or towards space (relatively ephemeral, but amenable to transportation). He argues that the “bias of communication” shapes the form of a society, and vice versa.

Edwards analyzes infrastructure in terms of scale, arguing that the analysis of infrastructure must proceed from a meso-scale, in between the micro-scale of everyday life, and the macro-scale of historical change. At a macro-scale, infrastructures can take decades to be constructed, and may be superseded by other infrastructures, or worn down by natural processes, over time. At a micro-scale, infrastructures provide the substrate over which everyday life proceeds. In Edwards’ argument, a meso-scale approach calls for attention to the historical development of infrastructures, and the institutional forms which develop, manage and regulate infrastructures.

Edwards’ arguments develop the approaches of the group of scholars involved in “large technical systems” research (Mayntz and Hughes 1988; Coutard 1999). These scholars focus on determining the historical stages of the development of infrastructure, building on the

classic theorization from Hughes (1983) of stages of invention and development, of technology transfer between societies and regions, of system growth, of the acquisition of momentum through technological and political-economic arrangements, and finally of the rise of planned regional systems. To a large extent, these scholars focus on technological form as it is developed and managed by state and corporate interests. A partial exception to this pattern may be found in Chatzis (1999), who reformulates Hughes' periodization to focus on the role of the community of engineers in developing standards and practices in the first stage, which are then taken up by corporations in subsequent stages.

More recently, the lessons from the large technical systems literature have been applied to the development of "cyberinfrastructure", the infrastructure which enables collaboration amongst networks of scientific researchers (Edwards et al. 2007). Research into cyberinfrastructure is distinctive from the large technical systems approach, in that it focuses not on infrastructures in relation to society at large, but rather on infrastructures developed and used for specific research contexts. Research into large technical systems and cyberinfrastructure holds several significant lessons (Jackson et al. 2007). First, that infrastructures cannot be developed entirely to a plan; they change through processes of scaling, technology transfer, consolidation and so on. Second, that initial choices - such as standards - can and do condition the shape of an infrastructure as it develops and grows. Third, that gateways for translating data between different local systems are of great importance. Fourth, that while tensions may restrict the process of infrastructure development, they are also productive sites for the determining choices towards the future evolution of infrastructure.

Star and Ruhleder (1996) apply the method of "infrastructural inversion" (Bowker 1994) to focus on infrastructural relations as the locus of change, rather than privileging people or things independently. They present several important properties of infrastructure. First, infrastructure is embedded in social and technological arrangements. Second, infrastructure is not visible in everyday use. Third, infrastructure has a spatial and temporal reach beyond a single site, or a single event. Fourth, infrastructure comes to be taken for granted through a process of learning to be a member of the community which uses the infrastructure. Fifth, infrastructural form and practices of users of infrastructure evolve alongside one another. Sixth, infrastructure is embodiment of standards. Seventh, infrastructure must be built on an installed base, which shapes the form that infrastructure takes. Eighth, infrastructure becomes visible only upon breakdown. Ninth, infrastructure is never changed through control from above, but rather develops through a variety of local changes.

Bowker and Star (2000) notably expand the concept of information infrastructure from problems of technology, to problems of classification. They show how categories, and the administrative structures which manage them - whether of disease or race - act as infrastructures over which society is organized across time and space.

I draw on these various theoretical frameworks to form an analytical model which calls for a focus on standards, institutions and political-economic arrangements. Star and Ruhleder (1996) deal in large part with the practices of users of infrastructure, separating users and designers. In contrast, I am concerned with the practices of those involved with operating and maintaining infrastructure, and the ways in which these practices overlap with the

practices of institutions, such as those which are responsible for standards. These practices are performed by engineers or administrators (similar to the engineers in Chatzis 1999), who are distinct from users or designers. I call this work “the practice of infrastructure”, to distinguish it from other senses of practice in relation to infrastructure. The properties of infrastructure that Star and Ruhleder (1996) present continue to hold true for my sense of practice: it is learned as part of a community of practice, and infrastructure does co-evolve with practice. The difference between my concept of practice and that elaborated in Star and Ruhleder (1996) is in the location of practice in operating and maintaining infrastructure, rather than in the use or design of infrastructure.

When the Internet is described as infrastructure, it is typically approached in one of two ways. First, in terms of features of the surface that ordinary users see, such as the Domain Name System (DNS), which maps human-readable names for websites, mail servers and other Internet services to Internet Protocol (IP) addresses (Sandvig 2013). Second, in terms of the materiality of the actual hardware involved in operating Internet infrastructure (Blanchette 2011). As I noted in Chapter 1, my focus lies in between these two positions: at the logical layer of the Internet, which mediates between the physical space of hardware, and virtual space as represented by domain names. This is not to say that physical hardware is not infrastructure, or that the DNS should not be thought of as infrastructure; rather, I mean to specify my use of the term “Internet infrastructure”. When I use this term, I am referring to the specific concept of infrastructure at the Internet’s logical layer, and its function in the production of virtual space over physical space.

2.2 Methods

2.2.1 Protocol as a Field Site

The objects of my research are a protocol - the Border Gateway Protocol (BGP) - and the infrastructure which it is used to construct, the Internet’s inter-domain routing system. How can a protocol be studied ethnographically? In this section, I briefly discuss the methodological concerns that arose as I sought to make sense of protocol as a fieldsite.

Protocols are standards which are used to construct the Internet. As such, they belong to the larger analytical category of standards involved in the construction of modern societies, from standardized voltages in power supplies, to standards for the purity of drinking water, to standardized layouts for computer keyboards. Standards shape society, just as society shapes standards; and standards provide the means through which social relations are stabilized over space and time through infrastructures (Bowker and Star 2000:13-14). Standards, and infrastructures, therefore provide a critical perspective on the production of space, following the analytics of place, territory, scale and networks.

Infrastructures, argue Star and Ruhleder (1996), are always relations, and never just things, a perspective which calls to attention the relations - both social and technical - which are involved in the production of infrastructure. They are embedded, “sunk into and inside

other structures, social arrangements and technologies” and learned through membership in a community of practice (Star 1999:381). The concept of practice at work here is one of *spatial practice*, since infrastructures are integral to the production of space. This spatial practice involves a spatial division of labor, albeit not one in which labor is specialized across different locations, but rather one in which similar labor is performed across space in order to construct a uniform infrastructure.

To understand the propagation of practice across space, a relational comparison must be made between the instantiation of practice in different locations. This requires a multi-sited ethnography, following people, things, and other analytical objects, across space and time (Marcus 1995). Yet this is not a picture of separated field sites, but rather of related field sites which may be conceived of relationally in the construction of infrastructure. The field site is no longer, then, a spatially bounded entity which can be entered into, and exited from; instead, it may be conceived of as a network form, in which a researcher’s task is to follow relations, and identify the critical places at which relations converge (Burrell 2009).

The performance of similar practices across space requires a disembedding of practice from its origin to imagine it as a global absolute, and a re-embedding in a remote location. My use of disembedding and re-embedding here is quite distinct from Giddens (1991). Giddens is concerned with the disembedding of social relations from local contexts to extend them across space and time, and their re-embedding in local contexts through interaction with abstract systems. In contrast, my conception of a disembedded practice is one which is *explicitly* coded in documents which describe that practice, and may be used for its instruction, and *implicitly* carried by the instructors who must be able to translate practice from its point of origin to a remote location. The process of re-embedding a practice does involve the transfer of knowledge - both explicit and implicit - but also involves a re-articulation of the practice to suit local conditions. Yet, since this practice involves a common protocol standard, the variations in local re-articulations of practice are limited. A protocol which defines the mechanisms for the interconnection of networks - such as BGP - requires certain minimum common conceptions of practice across all interconnected networks.

The practices involved in network interconnection are distributed to a large degree, just as the networks which they interconnect are distributed across the common topology of the inter-domain routing system. However, certain functions required for these practices are centralized; specifically, standardization and resource allocation. Centralized institutions are required to set common protocol standards, and allocate unique resources, both of which are required for the practice of inter-domain routing. Accordingly, I use the analytical triad of topology, standards and resources for my research.

In hindsight, all research is well planned. The practice of research, however, is a messy, sometimes unpredictable process, fraught with unexpected twists and turns. I originally planned only to study the practice of network administration involved in inter-domain routing. However, as my research progressed, I found that it was necessary to pay attention to the practices of standards setting and resource allocation, insofar as these were of importance to the practice of inter-domain routing. Similarly, I found that I had to turn my attention to the role of the state in India, even though this was not my original intention, as it became

apparent that the state was intimately involved in the production of the inter-domain routing system in India. These twists and turns may have been unplanned, but they represented critical inflexion points through which I was able to formulate the analytical framework of topology, standards and resources for understanding Internet infrastructure.

I used a large and varied corpus of materials for my research, from interviews and field notes, to publicly available archives of email lists, meeting attendance records, transcriptions and videos from meetings, and standards and policy documents, which I detail below. Inductive coding of these materials was useful to a degree for my analysis, to identify themes such as trust. However, I relied to a much larger extent upon thematic memos, which I developed throughout the course of my research, as these were much more useful for capturing new directions, and connections between different themes.

The overall research strategy which I adopt for this dissertation is a mix of interviews, participant observation, and document analysis. I interviewed network administrators responsible for inter-domain routing in North America and South Asia, to understand their practices of network administration. I approached interviewees through a combination of cold calls and snowball sampling through introductions provided by prior interviewees. My cold calls were targeted towards certain profiles: I aimed to get a spread of junior and senior network administrators, and network administrators who occupied certain roles within the network administration community. These roles varied from those of regular attendees at conferences (identified through analysis of conference attendee lists) to membership in the various committees governing the professional organization of the network administration community.

I performed additional interviews in India with bureaucrats in the Indian government, and with representatives of the ISP Association of India (ISPAI) as it became apparent that the Indian government and ISPAI were involved in processes which were provisioned by the network administration community, or by corporations, in North America.

I conducted fieldwork primarily at meetings of the North American Network Operators Group (NANOG) and the South Asia Network Operators Group (SANOG), which are professional organizations for network administrators in North America and South Asia respectively.¹

Together with the interviews I conducted, these field sites offered the opportunity to construct relational comparisons to make sense of the practices of network administration in the global inter-domain routing system. These two field sites present useful points of comparison, since North America is relatively central to the inter-domain routing system, while South Asia is relatively peripheral. Accordingly, they allowed me to make sense of how my analytics of topology, standards and resources are articulated differently in different contexts, while still being integrated into a global system.

¹There are numerous other network operators groups around the world, some operating at a regional level, and some at a country level. These include the Africa Network Operators Group (AfNOG), the Middle East Network Operators Group (MENOG), the Caribbean Network Operators Group (CARIBNOG), the Latin American and Caribbean Network Operators Group (LACNOG), the Japanese Network Operators Group (JANOG), the German Network Operators Group (DENOG), and many more.

To understand the practices involved in operating the centralized governance institutions of the Internet, I conducted fieldwork at meetings of the Internet Engineering Task Force (IETF), which is responsible for setting technical standards, and at meetings of the Internet Corporation for Assigned Names and Numbers (ICANN) and the American Registry for Internet Numbers (ARIN), which are responsible for managing critical Internet resources, such as Internet Protocol (IP) address space.

I studied email lists used by these communities and governance institutions to understand the nature of ongoing conversations which span individual meetings. I also studied various policy and standards documents to gain the background necessary to make sense of my fieldwork, and to understand the historical evolution of the inter-domain routing system.

This research was conducted with approval for human subjects research from the Committee for the Protection of Human Subjects at the University of California, Berkeley, protocol numbers 2009-4-30 and 2013-04-5171.

2.2.2 Meetings

I attended a range of meetings during the course of my research, from meetings of network administrators, to meetings of the organizations involved in standards setting and resource allocation. Meetings are named following the convention which is common amongst the organizations involved in managing Internet infrastructure: the organization name followed by the meeting number.

To study the practice of inter-domain routing, I attended meetings of the North American Network Operators Group (NANOG) and the South Asia Network Operators Group (SANOG), which are professional organizations of network administrators. I attended four NANOG meetings, and two SANOG meetings:

- NANOG49, held from June 13th to 16th, 2010, in San Francisco, California, USA.
- NANOG56, held from October 21st to 24th, 2012, in Dallas, Texas, USA.
- NANOG57, held from February 4th to 6th, 2013, in Orlando, Florida, USA.
- NANOG58, held from June 3rd to June 5th, 2013, in New Orleans, Louisiana, USA.
- SANOG18, held from September 8th to 16th, 2011, in Pokhara, Nepal.
- SANOG22, held from August 5th to 13th, 2013, in Mumbai, India.

I attended one full meeting of the American Registry for Internet Numbers (ARIN) - which manages Internet number resources in North America - and Public Policy Consultation (PPC) sessions held during NANOG meetings:

- ARIN30, held from October 24th to 26th, 2012, in Dallas, Texas, USA.
- ARIN PPC at NANOG57, held on February 5th, 2013, in Orlando, Florida, USA.

- ARIN PPC at NANOG58, held on June 4th, 2013, in New Orleans, Louisiana, USA.

I attended one meeting each of the Internet Engineering Task Force (IETF) - which sets Internet protocol standards - and the Internet Corporation for Assigned Names and Numbers (ICANN) - which is the global root authority for Internet resources:

- IETF74, held from March 22nd to 27th, 2009, in San Francisco, California, USA.
- ICANN40, held from March 13th to 18th, 2011, in San Francisco, California, USA.

I maintained extensive field notes for all of these meetings. Whenever I cite material from field notes in the chapters which follow, I identify the material using the meeting name, and the line number from the field note, prefixed by “F-”, in square brackets. For instance, [F-NANOG56:87] indicates material drawn from the 87th line of the field note for the NANOG56 meeting. I used attendance records, video of sessions, and transcripts - which are all publicly available on meeting websites - for my research, alongside my field notes. Whenever I draw from these secondary materials, I provide the link at which these materials are available on meeting websites.

Software for the analysis of meeting records was written in Python, and used a MySQL database for data storage. Graphing was done using the Python matplotlib library.²

2.2.3 Texts

I drew on a variety of publicly available textual materials for my research. In addition to the meeting transcripts listed earlier, these include, but are not limited to:

- Archives of the NANOG email list, available at <http://nanog.org/list/archives>.
- Archives of the SANOG email list, available at <http://news.gmane.org/gmane.org.operators.sanog/>.
- Archives of the ARIN Public Policy Mailing list, available at <http://lists.arin.net/pipermail/arin-ppml/>.
- ARIN policy documents, available at <https://www.arin.net/policy/index.html>.
- IETF Request For Comments (RFC) documents, available at <http://www.ietf.org/rfc.html>.

Whenever I draw from these textual materials, I either provide a full citation (for IETF RFCs) with a link in the citation, or a link in a footnote pointing to the location at which the material is available on the relevant website. All links listed here were last retrieved on July 23rd, 2014.

²Available at <http://matplotlib.org/>, last retrieved Jul 30, 2014.

2.2.4 Interviews

I conducted formal interviews with 57 interviewees, in addition to numerous informal conversations which I captured in my field notes. Formal interviews were, for the most part, recorded with oral consent from interviewees. In a small set of cases, interviewees were uncomfortable with recording, but consented to have me take notes from the interview. All recorded interviews were transcribed. Interviews were conducted over the phone, in interviewees' offices, or during the conferences I attended. Interviews lasted between half-an-hour to one-and-a-half hours. In some cases, I conducted follow-up interviews to clarify and expand on interviewees' earlier responses, or track changes over time.

I approached interviewees through a combination of cold calls and introductions facilitated by prior interviewees. While some cold calls were opportunistic, others were targeted based on a prospective interviewee's profile. For instance, I targeted interviewees based on their experience levels, the frequency with which they attended NANOG meetings, their participation on the NANOG email list, their role on SANOG committees, or their involvement with key organizations, amongst other things.

Interviewees were, for the most part, male. They varied in experience from 2 to over 30 years. The breakdown of interviewees by the sites of my research is as follows:

- Network administrators in North America, including employees of network equipment vendors: 40
- Network administrators in South Asia, including instructors traveling to South Asia from other parts of the world: 11
- Government officials, and policy advocates from the private sector in India: 6

Although South Asia appears to be under-represented, the lower number is indicative of a preference for casual conversations (which I captured in my field notes), rather than formal interviews.

All interviewees are quoted by a pseudonym, which is the interview number that I maintain in a spreadsheet. If I conducted a follow-up interview with an interviewee, I identify the original interview with a “-1” suffix, and the follow-up with a “-2” suffix. Whenever I reference an interview, I cite it with the interview number, and the page number of the interview transcription, or the line number if the interview was captured as notes, prefixed with an “I”, in square brackets. For instance, a quote from page 5 of interview 23 would be cited as [I23:5].

I archived personal email communications from my interviewees involving clarifications to points raised in interviews. With their permission, I have introduced issues raised in these emails to my analysis. Whenever I reference a personal email communication, I cite it with the keyword “EMAIL” followed by the date of the email, in square brackets. For instance, a quote from an email received on April 19th, 2013, would be referenced as [EMAIL:04-19-2013].

Chapter 3

A Social History of the Internet

In this chapter, I trace the development of the infrastructure and technologies of the Internet, with an explicit focus on aspects of the culture, institutions, and practices making up the social context in which these technologies were developed. While I separate social and technological elements for analytical purposes, it is important to note that the technological and the social constituted a unified reality for those who made the history of the Internet. In my historical account, social and technological elements in the early Internet were intertwined in their development. To disentangle these intertwined elements, and to provide a common analytical framework across this dissertation, I perform my analysis using the analytical triad of standards, resources and topology described in Chapter 2. As I will show, the topology of Internet infrastructure is not merely an implementation of technology standards and resource allocation, but rather has the capacity to influence standards development and resource allocation policy.

The mechanisms for the management of standards, resources and topology for network interconnection (and the Internet more generally) were formed in the collegial social context of computer science researchers, which, as I will show, had material implications for the technological and social forms of distributed governance in the Internet's inter-domain routing system. This chapter is as much an elaboration of technological concepts - such as packet switching, layering and the Internet Protocol - as it is of the historical context in which they were developed.

Through the socio-technical telling of history in this chapter, I illustrate the emergence of centralized institutions of power, and the distributed systems of practice and cooperation that stabilize the distributed computer network that is the Internet. While there are other histories which provide a detailed account of the development of the early Internet - such as Abbate (1999) from which I draw extensively - I distinguish mine through a particular focus on the technologies and social formations developed for the interconnection of networks. The Internet, after all, is not one network, but an instance of an *internetwork*: a network of many networks. Accordingly, the technologies and associated social formations required to manage the interconnection of networks are perhaps the most important components of the Internet.

This chapter sets the context for the remainder of this dissertation, in which I trace the

development of the culture, institutions and practices of the nascent Internet as they spread globally, transitioning from an academic research environment to the commercial Internet.

3.1 Sociability in the Early Internet

An oft-quoted origin story for the Internet is that it was created as a distributed system in response to Cold War era concerns of nuclear attack.¹ It was thought that a distributed system would better be able to survive a nuclear strike, as the surviving nodes of the system would be able to redirect messages around the failed nodes, “routing around failure” as the saying goes.² As with all origin stories, there is a kernel of truth to this claim; albeit a kernel surrounded by a substantial husk of extrapolation. It is important to get to an accurate version of this story since it sets up the social context within which the technologies of the Internet were developed.

The state-of-the-art in telecommunications at the time was the circuit-switched network. This was a system in which messages passed from one node to another through centralized switchboard. If the switchboard was disabled, for whatever reason, all nodes attached to it would become unavailable. Existing approaches to survival of communications in the event of a nuclear attack involved "hardening" these links to ensure that they would be able to physically survive a nuclear attack. In the early 1960s, Paul Baran at the RAND institute offered an alternative in the form of distributed system based on packet-switched networking. Rather than relying on a centralized switchboard, Baran’s packet-switched network broke messages up into small “packets” (or “message blocks” as Baran termed them) which would then be adaptively routed from sender to receiver, with each intermediate node making decisions about how to direct a packet next. A given message would be broken up into a series of packets, each of which might take different routes to reach its destination. The failure of a given switching station within a packet-switched network would not spell failure for the network as a whole; packets would be adaptively routed around the failed node to be delivered to their destination.

In the mid-1960s at the United Kingdom’s National Physical Laboratory, Donald Davies was grappling with another limitation of circuit-switched networks: if a circuit between a switchboard and a node was already in use, no-one else would be able to access that circuit. This was an important problem for time-sharing computers³, as phone lines would have to be kept open for each connected user. Davies developed his work independently with no knowledge of Baran’s earlier work; in fact, he came to know of Baran’s work only through a conversation with someone from the British Ministry of Defence who was familiar with Baran’s work via military contacts. Davies’ goal was to provide for multiplexed access to

¹For examples of this narrative, see Galloway (2005), Castells (2000:45), and even cybersecurity policy documents such as the US government’s 2009 Cyberspace Policy Review (United States Government 2009).

²While this saying was originally a technical claim, it has expanded to include the connotation that the Internet can route around political “failures” such as censorship.

³Mainframe computer systems on which processing time was shared amongst multiple users.

time-sharing machines, to which end he conceived of and built a system in which multiple data streams could be sent at the same time over the same circuit: unlike a circuit-switched network, a packet-switched network can alternate packets from different data streams over the same physical circuit.

Davies' work was in turn taken up by those pioneering the nascent ARPANET: a network funded by the US Department of Defence Advanced Research Projects Agency (ARPA). Computing resources were a major expense for ARPA, and it was thought that a network to allow the computing resources installed at particular research centers to be shared by geographically remote researchers would reduce ARPA's costs. ARPA's response was to fund the creation of a network - the ARPANET - which would provide connect these mainframe computers to research institutions. The researchers involved in creating the ARPANET were unaware of Baran's original work, but came to these ideas via a paper that a colleague of Davies' presented at a symposium in Gatlinburg, Tennessee, in 1967, taking on Davies' term "packet switching" for the technology they would apply to the ARPANET. Although the Davies' and Baran's approaches to the technology of packet-switching was similar, what is most important is the intention behind the application of this technology, which shaped the organizational and cultural setting of its development. In the case of the ARPANET, the intention was much closer to Davies' ideas of time-sharing than it was to Baran's ideas of survivability in the face of nuclear war (Hafner and Lyon 1998). In short order, a technology once intended to survive a nuclear attack found a much more mundane application in networking for access to time-sharing computer resources. As Abbate (1999:23) puts it, "If the watchword for Baran was survivability, the priority for Davies was interactive computing".

While the early development of Internet technologies may have their origins in the Cold War era, they were not entirely shaped by the political imperatives of that time. Packet-switched networking may have been developed with the military problem of survivability in mind, but it was eventually put to use for the fairly utilitarian end of access to time-sharing computer resources. Accordingly, the social context within which these technologies were developed was driven by academic researchers, rather than military priorities. The specific forms that these technologies took followed the contours of this social context. As Stephen Carr, Stephen Crocker, and Vinton Cerf - some of the key figures in the development of the ARPANET - relate:

We have found that, in the process of connecting machines and operating systems together, a great deal of rapport has been established between personnel at the various network node sites. The resulting mixture of ideas, discussions, disagreements, and resolutions has been highly refreshing and beneficial to all involved, and we regard the human interaction as a valuable by-product of the main effort. (Carr et al. 1970)

This collegial social context, and the imperatives attached to the development of the ARPANET, had material outcomes as they influenced the form of the technologies that were developed. The networks developed by Davies and the ARPANET team dropped certain

characteristics of interest to the military that Baran had privileged, such as high redundancy and cryptographic capabilities, while at the same time they took on some of Baran's ideas which fit their needs, such as adaptive transmission and efficient packet switching (Abbate 1999:39-40). Indeed, David Clark, another key figure in the development of the Internet, notes that accountability was the lowest ranked amongst the design priorities for the Internet protocols. The more highly ranked design priorities related to the interconnection of networks, the ability to continue service in the face of failures of individual networks or gateways, support for a variety of communications services and network architectures, distributed management of resources, cost effectiveness and low effort for the attachment of new hosts. Clark (1988) comments that "an entirely different network architecture would result if the order [of priorities] were changed."⁴

3.2 Building the ARPANET

The development of packet-switched networking was not without challenges; indeed, it met with much skepticism from the computer science research community, who believed that certain problems might prove insurmountable, such as the requirement to reorder packets which might arrive out of order at a receiving node. Another challenge arose from the need to interconnect a variety of machines from different manufacturers which were incompatible with one another. These and other challenges were addressed in the development of the ARPANET, which required a considerable effort in software and hardware redesign to implement standardized network interfaces on a large variety of computer architectures.

A key design principle adopted to address these issues was that of *layering*. In such a system, different layers take responsibility for particular functions, with well-defined interfaces to one another. Each layer can then be written to operate largely independently of other layers. In the case of the ARPANET, the problem of adapting a variety of computer architectures to ARPANET's standardized communication protocols was addressed by shifting the functionality for routing messages around the network to a specialized minicomputer, which was called the Interface Message Processor (IMP). As a result, computers connecting

⁴It is also important to consider that the development of these technologies were taking place against the backdrop of the counterculture movement. For instance, Turner (2006a) traces the connections between the "cyberculture" of Silicon Valley and the counterculture movement. While it is difficult to trace direct connections between the computer scientists involved in the development of the nascent Internet and the counterculture, it may be said with some certainty that these cultural threads became intertwined over time. By 1984, Stewart Brand was making claims about what information wants at the first Hackers Conference: "Information wants to be free. Information also wants to be expensive. ... That tension will not go away." In a similar vein at the same conference, arguments were taking place pitching ideas of free and open source software against those of commercial closed source software (Turner 2006b). The distributed network architecture enabled by Internet technologies played well with counterculture/cyberculture ideals of community and independence from central control. By 1996, these two cultural threads were so well integrated that Barlow (1996) could present "A Declaration of the Independence of Cyberspace" and Castells (2000) could theorize a "network society" which took on the characteristics of the computer networking technologies which enabled it.

to the ARPANET only needed to implement a standardized network interface to connect to their local IMP, which then took on responsibility for forwarding packets towards their destination. The IMPs were connected to one another, creating the actual packet-switched network. In effect, computing resources connected to IMPs (called “hosts”) formed a “host layer”, while the IMPs formed a “network layer”, with each layer operating fairly independently of the other. An additional “application layer” was designed above the host layer to provide specific kinds of application functionality such as file transfer between hosts.⁵

Lawrence Roberts, a computer scientist from MIT, was tasked by ARPA’s Information Processing Techniques Office (IPTO) with heading up the ARPANET effort. He was faced with the unenviable task of assembling a diverse group of researchers, contractors, and personnel managing existing university computing resources, to create a network based on an unproven technology. When necessary, he did wield ARPA’s financial clout to ensure cooperation from reluctant participants; but most of all, he was responsible for creating and maintaining the informal work practices and coordination for the basic research and development involved in the development of the ARPANET. This found form in the Network Working Group (NWG), created in 1969, an informal research group spanning several universities which included many graduate students who would go on to become key figures in the development of the Internet, such as Stephen Crocker, Vinton Cerf and Jon Postel. Roberts tasked the NWG with designing an important aspect of the ARPANET: the protocols required for the operation of the host layer. As Cerf (1990) said (originally quoted in Abbate 1999:73), “We were just rank amateurs, and we were expecting that some authority would finally come along and say ‘Here’s how we are going to do it.’ And nobody ever came along.”

The NWG adopted the convention of documenting their communications, proposals and meeting minutes in a series of documents which were termed “Request For Comments” (RFC). Setting norms for interaction, Crocker states in RFC 3 that the intention was to allow for the free flow of ideas, with minimal expectations of the content contained within an RFC: “The minimum length of a NWG note is one sentence” (Crocker 1969). The early RFCs covered a range of materials, including specifications of protocols (e.g., RFCs 1 and 2, specifying the host protocol), membership in the NWG which was “not closed” (RFC 3), notices of impending action (e.g., RFC 45, which notifies the arrival of a “clean version of the network protocol”), and the assignment of network resources (e.g., RFC 349 for the assignment of socket numbers⁶, or RFC 790 specifying assignments of numbers including sockets

⁵Although networks are commonly thought of as horizontal forms of organization, the vertical layering of the Internet illustrates how networks may form in independent layers which depend upon one another. Indeed, as I will discuss later, the Internet Protocol operates the critical “network layer” which forms the common network which unifies disparate physical telecommunications infrastructures, and provides a common platform over which disparate higher layer services may be constructed. The Border Gateway Protocol, which I discuss at length in Chapter 4, provides the means through which networks interconnect at the “network layer” on the modern Internet.

⁶Well-known socket (or port) numbers are used to identify a channel over which a specific protocol (such as file transfer or terminal access) operates.

and Internet Protocol address space). Jon Postel took on two important administrative roles beginning with the NWG, acting as the editor for the RFC series, and as the gatekeeper assuring uniqueness of assigned network resources (such as well-known port numbers and IP address space); these were both roles that he would continue to be involved with until the time of his death in 1998.⁷

While some ARPANET development was outsourced to researchers, such as the NWG, other pieces of the ARPANET were developed by private contractors. Notable amongst these was the firm of Bolt, Beranek and Newman (BBN), which was responsible for developing the IMP hardware and software, and maintaining the network of interconnections amongst IMPs. The layering of the ARPANET network architecture allowed for separation of administration of layers, just as it allowed for separation of development activities for each layer. Notably, this (vertical) separation of administration of layers also implied a (horizontal) separation of administration between individual connected sites within the ARPANET. The administrators at any computing facility connected to the ARPANET could continue to operate their systems as they saw fit, with their only investment being in the implementation of the host protocol to connect their computer systems to the local IMP. These arrangements were the genesis of more formalized mechanisms to maintain the autonomy of administrative domains which emerged as the Internet evolved. The ARPANET contract called for an initial connection of 4 sites which were ARPA-funded computing facilities: the University of Utah, the Stanford Research Institute in Menlo Park, the University of California, Los Angeles, and the University of California, Santa Barbara. Once the ARPANET concept was proven in this implementation - on December 5th, 1969 - it was expanded to cover other sites as well. The ARPANET expanded rapidly after this initial implementation, growing to 37 sites in 1973, including international links to Norway and London. By this time, the ARPANET had also transformed from being an experimental system in which failure could be tolerated to being an essential tool for its users. For instance, the link to Norway was put in place to connect a seismic monitoring facility which was used to monitor (amongst other things) Soviet nuclear tests; where analysis of the data used to take weeks or months before (as data had to be shipped on tapes via courier), the new system reduced this delay to mere days. Given this shift from experimental to production network used on a daily basis by research and military sites, the administrative responsibility for the ARPANET was shifted from ARPA to the Defence Communications Agency (DCA) in 1975. As an agency focused on research and development, ARPA was ill-equipped to administer a production network; the DCA, in contrast, was tasked with performing such functions for the military. By 1983, a separate military network, MILNET, was split off from the civilian ARPANET⁸.

As this account illustrates, while the ARPANET was not developed for military pur-

⁷Postel proposed that he become the “czar” for socket numbers in RFC 349. The assignment of network resources cannot be automated since it involves the administrative work of uniquely assigning resources to interconnected networks administered by different entities, such as different research institutions and universities. I discuss this issue further, as it relates to IP address space, in Chapters 4 and 5.

⁸MILNET eventually became the Non-classified Internet Protocol Router Network (NIPRNet), which - as the name suggests - is used for non-classified network traffic by the US military.

poses, it certainly found some military applications early in its history. Similarly, while the ARPANET was primarily a research enterprise, it did depend on private contractors to build and operate certain of its elements. The development of the ARPANET was not by any means an idealized research endeavour; it was enmeshed in economic and political concerns through and through. It fell upon the ARPANET managers, led by Lawrence Roberts, to ensure that these multiple, often conflicting, interests were reconciled and re-formulated so as to insulate researchers - such as the NWG - and allow them to develop technologies with engineering priorities in mind, in a relatively informal cultural environment. This account also shows how a separation of concerns was constructed amongst elements of network topology (implemented by BBN), standards development (the NWG's RFC process), and the allocation of network resources (Postel's management of socket numbers and IP address space): in part through technical decisions such as a layered network architecture, and in part through associated management strategies. These associated technological and social formations had a significant impact on the culture, institutions and practices of the Internet as it would develop in later years, allowing those involved in the development of the Internet to imagine their activities as being in service of "the good of the Internet",⁹ finding ways to reconcile their own political and economic concerns with this unifying goal.

3.3 The TCP/IP Protocol Suite

The ARPANET IMPs communicated with one another using a piece of software called the Network Control Program (NCP). One of the key assumptions underlying the design of the ARPANET was that one - and only one - host computer would be connected to any given IMP. The IMP was merely the host computer's channel to the ARPANET. The concept of a local area network (LAN), and even the term "Internet"¹⁰, had yet to be developed. The push for the technologies that would help build the Internet came as it became apparent that the ARPANET would not be able to deal with the data transmission mechanisms with different physical characteristics. The ARPANET backbone network, operated by BBN, was engineered to be highly reliable, and this assumption was built into the NCP, which assumed the presence of a highly reliable physical transport for data. As experiments with packet-switched technologies over terrestrial radio and satellite demonstrated, the reliability of the physical transport cannot always be assumed. In addition, different physical transports may have radically different transmission characteristics: for instance, satellite systems provide very high bandwidth, but also a very high end-to-end delay in the transmission of data, as the data travels from earth to the satellite and back again. In comparison, a wired network such as the ARPANET backbone, might provide a somewhat lower bandwidth, but a substantially lower end-to-end delay. Building a reliable data transport over any individual physical transport presents its own unique problems; building a reliable data

⁹A sentiment which I explore further in the following chapters.

¹⁰As late as 1982, the ARPANET was referred to as a "catenet", short for "concatenated network", in RFC 827.

transport across *different* physical transports is substantially more difficult. These challenges drove the creation of the Internet protocols which would eventually supplant the NCP.

The Internet protocols were first developed as the Transmission Control Program (TCP, eventually called the Transmission Control Protocol) in 1973, and documented in RFC 675 in 1974 (Cerf et al. 1974). As discussions about TCP progressed, it became apparent that TCP itself was being called upon to do too much work - and so, TCP was broken up into two layers, TCP and IP (Internet Protocol). IP is responsible for maintaining point-to-point links within an internet, while TCP assures reliable delivery of data across any two hosts in a network, potentially traversing multiple point-to-point links. Together, these core Internet technologies are known as the TCP/IP protocol suite.

TCP/IP introduced several innovations, not the least of which is a flat address space. Every machine on a TCP/IP network is assigned an IP address which identifies it uniquely; when the scope of this uniqueness is an internet, globally unique IP addresses present a critical mechanism for identifying endpoints for communication, as well as intermediaries involved in routing communications between endpoints. Importantly, at the TCP layer, the services involved in communicating with one another at endpoints can deal with each other directly, with no concern for how data will be transported between endpoints. It falls upon the IP layer to actually transport data between these endpoints, potentially across multiple point-to-point links.

IP addresses are 32 bits in length, offering 2^{32} possible addresses, and are typically written as four 8-bit segments (or octets) delimited by “.” marks. For instance, the web server providing access to the UC Berkeley website has the IP address 169.229.216.200. While the address space for IP is flat in theory, it formed a two-level hierarchy in its initial implementation. Where the NCP assumed that only one host connected per IMP, TCP/IP assumed that multiple hosts could be connected behind a gateway connecting them to the Internet. Apart from the problem of identifying individual hosts by IP address, there also emerges the problem of identifying gateways in order to route data traffic to the correct gateway which could then forward it on to the host it is intended for. The initial implementation of TCP/IP solved this problem by splitting IP addresses into two components: a prefix which identifies the network (i.e., the gateway) which the IP address belongs to, and a suffix which identifies an individual host within a network. Three classes of IP addresses were created with this convention in mind: class A addresses with an 8-bit prefix (with 2^{24} possible hosts), class B addresses with a 16-bit prefix (with 2^{16} possible hosts), and class C addresses with a 24-bit prefix (with 2^8 possible hosts) (Postel 1981). A special class D address space was introduced later, to provide multicast services (for one-to-many communications). Eventually, as worries about IP address space scarcity first surfaced in the 1990s, this system was generalized to allow for variable prefix lengths. Under this system, called Classless Inter-Domain Routing (CIDR), blocks of IP addresses are identified by a prefix and a number indicating the length of the prefix. For example, UC Berkeley has several IP address block assignments, one of which specifies addresses in the range 128.32.0.0-128.32.255.255. This could be viewed as a class B address with the network prefix 128.32.0.0, or it could be written as 128.32.0.0/16 in CIDR notation, indicating that the first 16 bits of the address (the first two octets) are

the IP address prefix, while the remaining 16 bits may be used for specifying individual IP addresses at UC Berkeley. In general, these allocations are termed as IP address prefixes, although the older term “networks” (reflecting the notion of a prefix being a network identifier in address space classified as class A, B or C) is also sometimes used. I will use the term “IP address prefix” to avoid confusion.

The IP addressing scheme creates a system in which identifiers are abstracted away from physical hardware and links, and in which administrators at individual networks are free to attach new hosts as they see fit, as long as they have sufficient IP address space within their allocations. The logic underlying IP addressing reflects two important themes. First, a separation from physical space through the abstraction of IP addressing. Second, the principle of autonomy, as administrators were permitted to construct their own networks internally as they wished, subject only to the constraints of a common set of network protocols (the TCP/IP protocol suite), and a common authority for the allocation of IP address space.

As I noted earlier, the task of ensuring uniqueness of allocation of network numbers - including IP addresses - was taken on by Jon Postel. One of my interviewees, who was responsible for a campus network at the time, commented:

There weren't that many people, there weren't that many requests. You sent a request to Jon, Jon sent you back an address. I heard tales, though I can't confirm they're true, that sometimes people ran into Jon in meetings, and he gave them some address space, wrote it down in a notebook, and hopefully put them in a database someplace, it appears he was good at that, Jon was good at a great number of things. His loss hurt the Internet a lot.¹¹ It was a very informal process in those days. [I17:28]

Where NCP - and the ARPANET - assumed direct links between IMPs, mediated by the BBN-operated network layer, TCP/IP enabled a vision of a more complex network topology, potentially spanning multiple physical media, and multiple IP routers, for data transport between two hosts. This was first demonstrated in 1977 when an experimental TCP/IP network was set up to show how the ARPANET could be connected to satellite and terrestrial radio networks (Abbate 1999:131). ARPANET eventually transitioned from NCP to TCP/IP in 1983 (Abbate 1999:142). This was, by all accounts, a painful process: switching from one network standard to another requires that all interconnected sites transition at the same time; otherwise, the network as a whole is effectively partitioned between those who use competing standards. The “flag day” for the switchover was January 1st, 1983, and many sites were not ready for the transition, causing disruptions in ARPANET access. Some of those involved handed out buttons labeled “I survived the TCP transition”¹².

¹¹Jon Postel passed away in 1998, and is remembered through a variety of means, including scholarships, a research center at USC (<http://www.postel.org>, last retrieved Dec 9th 2013) and RFC 2468, “I Remember IANA” authored by Vint Cerf.

¹²For instance, see this post to the Internet history email list: <http://www.postel.org/pipermail/internet-history/2009-April/000799.html>, last retrieved Dec 9th 2013.

With the possibilities for a more complex network topology enabled by TCP/IP came an additional challenge of routing data. NCP simply transferred data from one IMP to another; the Internet protocols would have to provide mechanisms to discover routes between and within networks. A variety of routing protocols were developed for this purpose, broadly divided into “interior gateway protocols” which provide for the discovery of routes within a network (behind a gateway), and “exterior gateway protocols” which provide for the discovery of routes between networks (between gateways). The Border Gateway Protocol (BGP), which is the subject of this dissertation, is an example of an exterior gateway protocol. In fact, it is the only exterior gateway protocol in use today. While networks may use whatever interior gateway protocols they wish within their boundaries, they must use a common exterior gateway protocol to form an interconnected group of networks, just as they must all conform to the same TCP/IP protocol standards, and rely on a common numbering authority. I will examine the development of exterior gateway protocols in detail in the sections that follow, showing how the forms of these protocols developed alongside particular network topologies and social contexts.

3.4 Building the NSFNET

It is notable that while the ARPANET served as a proving ground for packet-switched networking, it was by no means an internet (a network of networks); it was a significant achievement, integrating disparate computer architectures and operating over disparate physical media (radio, satellite and copper), but it formed a *single* multi-layered network. The TCP/IP protocols were developed in the context of the ARPANET, but it would be another decade, and separate development effort, before a true internet of interconnected, independently administered networks was constructed. This was in no small part due to the fact that as the ARPANET grew, it rapidly became a day-to-day utility, rather than a platform for research into computer networking technologies. While technologies may change and develop rapidly, the pace of change in infrastructures built with these technologies is often glacial in comparison, due to lock-in to widely deployed technology standards, and because of user expectations of reliability¹³.

The analytics set up in the previous sections show how the diversity, structure and size of the topology for an infrastructure has material effects on the infrastructure’s capability for change. As I shall illustrate in this section, the structural form of an infrastructure’s topology also interacts with efforts for standards development and resource allocation policy for an infrastructure. Topology is not merely implemented to follow standards and resource policy, but rather has the capacity to influence the development of standards and the setting of resource policy. The conceptual triad of topology, standards and resources becomes a critical lens through which to understand the development and governance of infrastructure. In

¹³For a more modern example of glacial rates of change in Internet infrastructure, we need only look to the years of effort it is taking to migrate the Internet from the current version of the Internet Protocol, IPv4, to the newer version, IPv6.

addition, the cultural and political-economic interests at play in the development of this triad must also be taken into account: for instance, the involvement of the military and private industry in the development of the ARPANET, and the collegial social context amongst the researchers involved with the ARPANET.

A true internet finally came into being with an NSF-funded effort to create a new packet-switched network to interconnect research institutions: the NSFNET. The initial intention behind the NSFNET - like the ARPANET - was to link six supercomputer centers (5 of which were NSF-funded). This early version of the NSFNET was created in 1985, running on 56 kbps links. Demand for access to the NSFNET increased rapidly, just as usage of the network quickly outstripped the capacity of these links. In consequence, the NSFNET went through a series of upgrades, eventually re-engineering its backbone to use T1 (offering 1.544 Mbps) and then T3 (offering 44.736 Mbps) links. Simultaneously, the NSFNET developed into a three-tiered hierarchical topology in which the backbone network connected to regional networks which in turn provided service to institutional networks.¹⁴ For example, UC Berkeley connected to BARRNET (the Bay Area Regional Research Network) which in turn provided connectivity to the NSFNET backbone. By 1988, the NSFNET connected thirteen regional networks and supercomputer centers, providing service to 170 campus networks¹⁵. As with the ARPANET, the NSFNET effort involved private industry as well as academic research efforts. Merit Network, based out of the University of Michigan, led a consortium including IBM (for network hardware) and MCI (for backbone links) to upgrade the NSFNET backbone. As the complexity and capacity of the backbone increased, a new non-profit entity - Advanced Network Services (ANS) - was created by Merit in 1990 to manage the NSFNET backbone, alongside a new for-profit entity to provide commercial Internet services (Harris and Gerich 1996).¹⁶

The Internet Engineering Task Force (IETF) was the site for coordination amongst all these various entities: campus networks, regional networks, the backbone network, network hardware vendors. It was to the NSFNET what the NWG was to to the ARPANET. The IETF was constituted as an activity of the Internet Activities Board (IAB, later called the Internet Architecture Board) in 1983, although it would be 1986 before the first IETF meeting was organized. The IAB was originally constituted as the Internet Configuration and Control Board (ICCB) by Vint Cerf in 1979, with the intent of providing wider research input into the ARPANET development process. In this period, members of the ICCB, and the later IAB, were selected by the IAB chairman, as a “council of elders” for the networking community (Krol and Hoffman 1993:4). As the NSFNET expanded the IAB was institutionalized into a more formal body with clearly articulated processes for the selection of its membership (Chapin 1992). In contrast, the IETF continued to function as the NWG

¹⁴Note that this kind of topology is referred to as a “tree” in computer science. I use the concept of “hierarchy” intentionally to show how administrative authority follows the hierarchy of networks in this topology.

¹⁵<http://www.nsf.gov/about/history/nsf0050/internet/launch.htm>, last retrieved Dec 9th, 2013

¹⁶Since NSFNET policies forbade commercial activity (following the NSF’s research mission), a separate entity had to be created to provide for the increasing commercial demand for access to Internet services.

did, with little notion of formal membership beyond participation in the IETF process, an ideal it continues to hold to this day. The IETF also continued the NWG practice of publishing RFCs covering a range of subjects, from standards to resource allocation, and provided a space for coordination amongst network administrators and researchers at the many networks and organizations making up the NSFNET. A formal organizational home for the IETF and IAB was constituted in 1992 as the Internet Society (ISOC), which engages in a range of activities, from participation in Internet policy debates to supporting the development of Internet expertise and shared infrastructure around the world.

As one of my interviewees, who was involved with managing the NSFNET backbone at Merit (and later ANS), noted, “People were very cooperative, we’re all thinking that we’re doing a great thing here, it’s historical, people just want to do good” [I5:5]. While I do not mean to suggest with this quote that the NSFNET was without conflict, it is notable that those that I’ve spoken to about it remember the NSFNET fondly, commenting on how much smaller and more tightly-knit the technical community was at that time, operating with comparatively little interference from political and market interests. Since the NSFNET was itself a research activity, and funded by the NSF to interconnect research institutions, it was in some ways even easier than it had been in the ARPANET to maintain and develop the network in a relatively informal social context, especially since many of those involved with the ARPANET carried over their ideals and practices to the NSFNET. What is remarkable is that this informality was maintained even as certain functions required for coordinating NSFNET activity became increasingly institutionalized, such as the IAB, the IETF and the IANA.¹⁷

These informal institutional contexts were elevated to more clearly articulated principles of organization and practice when the IETF networking standards came into conflict with the Open Systems Interconnection (OSI) networking standards specified by the International Organization for Standardization (ISO). Since 1977, the ISO had been developing its own networking standard, OSI, and it was thought that OSI would eventually supplant TCP/IP as an international standard for computer networking.¹⁸ Where TCP/IP was developed in response to particular problems in computer networking (such as building networks to span disparate physical media for data transport), OSI was created as a means to integrate the vendor-specific proprietary computer networking architectures of the 1970s. By 1990, the US government had standardized a version of OSI for government procurement purposes called the Government OSI Profile (GOSIP), published as Federal Information Processing standard 146 (FIPS-146).¹⁹ In consequence, the IAB set a direction for the IETF to create standards for interoperability between OSI and the Internet protocols (Cerf and Mills 1990).

¹⁷The Internet Assigned Numbers Authority, the formal title for the function that Jon Postel originally took on as the “czar” of numbers.

¹⁸Even though the OSI standards did not eventually succeed, their influence is so pervasive that the OSI seven layer networking model (in comparison, TCP/IP presents a five layer model) can be found to this day in computer science textbooks.

¹⁹Other governments, including the UK, New Zealand and Australia also created GOSIP standards of their own around this time. The UK standard evolved into a European GOSIP standard.

However, there were tensions between the TCP/IP camp and the OSI camp well before the IAB's push for interoperability. Those involved with the IETF viewed the OSI model as being overly prescriptive, comparing the OSI model to theology and mysticism. These tensions came to a head when the IAB formulated a proposal for replacing IP with an OSI protocol called CLNP (ConnectionLess Network Protocol), anticipating that the concern that IP was running out of address space for assignment might be mitigated by the use of CLNP, which offered a substantially larger address space. Given the tensions at play, the IAB intended to carefully engage in conversation with the IETF over this proposal; however, the IAB's proposal was reported on in the news before they could present it to the IETF themselves. Many at the IETF interpreted the IAB's proposal as selling out to OSI and engaging in top-down decision making which was contrary to the spirit of the Internet standards process, resulting in severely acrimonious relationships between the IAB and the IETF. The IAB was forced to confront this conflict at the July 1992 IETF meeting in Cambridge, Massachusetts. Vint Cerf removed his trademark three-piece suit on stage to reveal a t-shirt which stated "IP on Everything". In his plenary, David Clark (who chaired the IAB through the 1980s), coined a phrase which has since been enshrined amongst the IETF community as capturing the spirit and values of the Internet standards process: "We reject: kings, presidents and voting. We believe in: rough consensus and running code." (Russell 2006). Unlike the ISO, which requires formal votes by ISO members to ratify standards, the IETF relies on demonstrable consensus and working implementations to elevate a technology proposal to an Internet standard.²⁰

Even as those involved in the management of standards, resources and topology separated - creating formal institutional arrangements for standards development (the IETF) and for resource allocation (the IANA) - the culture of openness, transparency and cooperation originating in the context of the ARPANET was kept alive. I will show in the following section how this culture had a material impact on the development of exterior gateway protocols.

3.5 Networking the Networks

The first exterior gateway protocol was developed in this increasingly institutionalized, yet still loosely structured social context. Notably, it was developed to fit the hierarchical network topology of the ARPANET, in which a backbone network provided connectivity to nodes, which in turn provided connectivity to computers within the organizations hosting the nodes. The topology of the NSFNET followed a similar hierarchical pattern, with the NSFNET backbone network providing connectivity across geographies (within the USA and internationally), connecting to regional networks, which in turn provided connectivity to networks at research and education institutions. This hierarchical topological arrangement had significant implications for the form of the exterior gateway protocols which were developed in the context of the NSFNET, as I will show here.

²⁰I will engage more fully with the IETF process in the Chapter 5.

The first exterior gateway routing protocol was called, quite literally, the Exterior Gateway Protocol (EGP), originally specified towards the end of the ARPANET period in 1982 as RFC 827 (Rosen 1982). GGP assumed a full mesh of interconnections between gateways, which created significant difficulties in the administration of the ARPANET routing infrastructure, elaborated in RFC 827. The GGP routing algorithm required increasing overheads, as all gateways had to be in constant communication with all other gateways to keep routing information up-to-date. In addition, fault isolation, and changes to the routing infrastructure required coordination and collaboration amongst network administrators at all ARPANET gateways. As these difficulties illustrate, even in the relatively constrained context of the ARPANET, the interconnection of networks was simultaneously involved technical problems of efficiency and reliability, and social problems of coordination and collaboration.

EGP addressed the technical difficulties with GGP by introducing support for a hierarchical routing topology: rather than sending routing updates to every other gateway (as in the full mesh topology of GGP), gateways would only have to send routing updates to their parent and child gateways in the routing hierarchy. This significantly reduced the load on individual gateways, and increased the scalability of the EGP network topology.

EGP addressed the social difficulties with GGP by recognizing that routing is an inherently social problem, since gateways are operated by independent organizations. EGP specifies the concept of independently administered domains, or “autonomous systems”, which are interconnected by gateways. Each of these autonomous systems is free to operate internally as they see fit, both from an administrative and technical perspective. For instance, a network administrator for an autonomous system can choose an interior gateway protocol and internal network topology suitable to its particular requirement, with no need to coordinate with personnel outside of that autonomous system. It is notable that EGP assumes a hierarchical network topology amongst autonomous systems, reflecting the topology of the ARPANET, and later, the NSFNET:

It must also be clearly understood that the Exterior Gateway Protocol is NOT intended to provide information which could be used as input to a completely general area or hierarchical routing algorithm. It is intended for a set of autonomous systems which are connected in a tree, with no cycles. It does not enable the passing of sufficient information to prevent routing loops if cycles in the topology do exist (Rosen 1982:6).

The design of EGP assumes a particular structural model for network topology, in turn constraining the set of possibilities for governance and organization of an internet using EGP as its inter-domain routing protocol. A hierarchical network topology drives assumptions of centralized - rather than distributed - institutional arrangements, which would seem to run counter to the ideals and design goals of the Internet community. The assumptions built into the design of EGP must be understood as a pragmatic consequence of the topological arrangements of the NSFNET. Apart from the dependency on a hierarchical network topology, EGP had another important problem: all EGP gateways were expected to send a notice

of the IP address prefixes that they could reach to their neighbors every three minutes. This caused a significant load on EGP gateways, especially as the number of autonomous systems and IP address prefixes on the NSFNET increased.

In response to these limitations, a successor to EGP was specified in RFC 1771 (Lougheed and Rekhter 1989): the Border Gateway Protocol (BGP). BGP is sometimes called the “three napkins protocol”, because it was originally sketched out over lunch on three napkins by its authors, as Yakov Rekhter recalls in the preface to a textbook on BGP (White et al. 2004). BGP took on many of the concepts in EGP, such as the idea of autonomous systems and gateways, but made two crucial changes. First, a BGP router would be required to send a routing update to its neighbors only when there was a change to its routing information base²¹. Second, the routing announcements that BGP routers send to one another would include the list of autonomous system numbers through which the announcement had been relayed²². The first change had the effect of radically decreasing the load on BGP routers, as compared to EGP gateways, which was the immediate goal of BGP. The second change enabled more complex non-hierarchical network topologies, as it eliminated the possibility of routing loops; BGP routers would be able to discard routing announcements from neighbors which had already passed through them, since they would see their own autonomous system number in the routing announcement.

Routing announcements in BGP exhibit a curious property: there is no mechanism built into BGP to allow recipients of a routing announcement to establish its veracity. Any given autonomous system might make claims of IP address prefixes to which it has routes to its neighboring autonomous systems, but these neighbors have way to establish the validity of these claims.

Through my interviews, it became apparent that this was a consequence of strong relationships of trust amongst the NSFNET technical community. As an interviewee who worked on the NSFNET put it, “the security issue was one that was not addressed initially, because it was not a commercial Internet, everybody trusted everybody, we were all hackers” [I17:3].

This apparent lack of security was also a consequence of the hierarchical topology of the NSFNET: regional networks filtered announcements from attached campus networks to ensure that attached networks were only announcing address space which belonged to them, and the NSFNET backbone performed a similar function with regional networks. These functions were performed at the regional network and backbone network using a centralized database called the Policy Routing Database (PRDB). The social relationships needed for the administration of this database followed the links of the network topology - as an administrator responsible for the NSFNET backbone told me:

²¹Such as a new block of IP addresses being assigned to its network, or a loss of connectivity to a neighbor.

²²In the NSFNET, for example, a campus network such as UC Berkeley might announce its address space via BGP to its regional network, BARRNET, which would in turn relay this routing announcement to the NSFNET backbone, which would relay the announcement to other regional networks, which would then send the announcement on to their attached campus networks. Each intermediate network - or “autonomous system” - would append its autonomous system number to the BGP routing announcement as it relayed it onwards.

What we did was a lot of manual work. We had regional designated persons, we have certain names... so they submit a list of routes that we, NSFNET, would... they send us via email, OK, here's the list. Then we update the database, and we use the database to generate the configuration file, the accept list, you know, that we accept these networks. We updated it once a day, so at that time it seemed to be acceptable, people don't do things so quick, like these days. So, you know, people send email during the day time, and at night we update the database. ... So that's how we addressed the issue, we have a trusted source, which is a representative of the regional network, who tells us what prefix to accept. So if they don't tell us, and they announce a false prefix, we're not going to accept it, it's not going to do any damage to the integrity of the routing table. That's how things were operating up to 1995. [I5:6]

In essence, the problem of establishing the veracity of a BGP announcement was solved through - and a consequence of - a hierarchical network topology, which enabled the maintenance of a centralized routing information database, maintained through relationships of trust following the topology of the NSFNET. These trust relationships also represent social connectivities through which the practice of managing BGP was made sense of, and transmitted. The NSFNET backbone administrator I interviewed told me how she was responsible for helping many of those who were "trusted source" representatives of regional networks to learn how to configure and operate BGP routers. Similarly, a network administrator from a campus network told me how he, in turn, learned how to configure connections between the campus' BGP routers and the regional network's BGP routers from the representative of the regional network through which the campus connected to the NSFNET. It is difficult to disentangle the nature of BGP, the hierarchical topology the NSFNET in which it was first deployed, centralized facilities storing records of allocated resources, and the tightly-knit technical community who developed and managed this infrastructure. All of these elements fed into one another, each constantly developing in relation to the others, to construct a unified socio-technical system.

3.6 From NSFNET to Internet

The NSFNET was privatized to create the commercial Internet on April 30th, 1995. Commercial services were offered by NSFNET-related organizations prior to this date, but the acceptable use policy for the NSFNET forbade commercial traffic over the NSFNET backbone, greatly limiting the reach of nascent commercial Internet providers. Following the privatization of the NSFNET, a range of commercial providers in the USA invested in backbone network infrastructure of their own, spanning North America, and establishing Internet connectivity to international locations.²³ To put this complexity in perspective, the Internet

²³Regional providers of commercial Internet services existed both in the USA and the rest of the world prior to 1995, but it was only after the NSFNET was privatized that all of these disparate providers interconnected

is currently composed of over 45,000 autonomous systems, announcing more than 485,000 IP address prefixes.²⁴ This is a significantly larger and more complex topology than that of the NSFNET, better characterized as graph with no single point of control, rather than a hierarchy with the possibility of absolute control by the top level of the hierarchy.

BGP underwent a series of changes leading to BGP version 4 (Rekhter and Li 1995) at the time of the NSFNET's privatization. Numerous extensions have been added to BGP since then, but the fundamental "trusting" nature of the protocol remains.²⁵ If, as I have argued, the stability of the inter-domain routing system operating over BGP on the NSFNET depended upon tightly-knit social relations of trust and a hierarchical network topology, how would the inter-domain routing system be stabilized over a far more complex network topology, with multiple autonomous systems competing with one another at a variety of geographic scales?

The NSF anticipated these issues through strategies which aimed at preserving elements of the environment under which BGP operated in the NSFNET. While it would no longer be possible to maintain centralized control of routing via a hierarchical network topology, a centrally managed routing information database (like PRDB) could still be a viable measure, as could support for activities aimed at fostering coordination and relationships amongst technical personnel at the many networks that made up the nascent Internet.

The NSF addressed the first of these as part of a solicitation issued in 1993 for building a National Research and Education Network (NREN) for the provision of Internet service to academic and research institutions once the NSFNET was privatized (NSF 1993). This solicitation solicited proposals for three distinct components: Network Access Points (NAPs) at which networks could interconnect with one another, very high speed Backbone Network Services (vBNS) to provide Internet connectivity between academic and research institutions, and a Routing Arbiter (RA) to act as a centralized store for routing information. The call for the Routing Arbiter was quite clearly modeled on experiences from the NSFNET and PRDB:

The solicitation also invites proposals for an RA organization to establish and maintain databases and routing services which may be used by attached networks to obtain routing information (such as network topology, policy, and interconnection information) with which to construct routing tables. This component of the architecture will provide for an unbiased routing scheme which will be available (but not mandatory) for all attached networks. The RA will also promote routing stability and manageability, and advance routing technology. (NSF 1993)

to form a truly global Internet.

²⁴See the CIDR Report for up-to-date numbers: <http://www.cidr-report.org/as2.0/>, last retrieved Dec 11th 2013.

²⁵The IETF has developed a family of standards for securing BGP called Resource Public Key Infrastructure (RPKI), but this is not widely deployed at this time. I will discuss RPKI further in following chapters.

The RA faced a major challenge: unlike PRDB which was used to enforce policies due to its topological location at the NSFNET backbone network, the RA was to be a third-party service which “*may* be used by attached networks”; it did not *have* to be used, nor did it have any means to enforce the policies maintained in its routing information database. The contract for building and maintaining the RA was awarded early in 1994 to Merit Networks and the University of Southern California’s Information Sciences Institute. They created the Routing Assets Database (RADb)²⁶, which remains in use today; a shared, publicly available, store of routing information, with this information maintained voluntarily by network administrators from autonomous systems across the Internet. Several other parallel efforts developed as a globally coordinated set of routing information databases, collectively known as the Internet Route Registries (IRRs).

RADb remains, by far, the largest of these registries, in part because it was the first of the IRRs, but even more so because the organization that took over the NSFNET backbone, ANS, mandated the use of RADb to networks that wished to carry traffic across their backbone. ANS absorbed key personnel from the NSFNET backbone in addition to the infrastructure of the backbone itself, and as such, these administrators brought their practices to bear in the management of the new ANS backbone. ANS was one of the largest networks at the time, which made its policy of requiring the maintenance of routing information in RADb of material importance to almost every network attached to the nascent Internet, since most Internet traffic had to traverse infrastructure operated by ANS, and was therefore subject to ANS’ policies. As one of my interviewees commented, “If you didn’t put your data in RADb, you didn’t get routed by ANS, which was unacceptable” [I17:4]. ANS was able to leverage its topological position as a point of control which had to be traversed to force other networks to abide by its policies. This changed, of course, as other backbone networks were developed to compete with that operated by ANS, removing the possibility for any single point of control in the topology of the inter-domain routing system.

There were several problems with RADb, not the least of which was that it depended upon good housekeeping by network administrators at hundreds - and then thousands - of networks across the Internet. This problem was magnified manifold as network administrators switched organizations and as mergers and acquisitions swept through the ISP industry in the late 1990s. In addition, in the initial implementation of RADb, it was difficult to establish the latest updates to the database, certain validations for data entry weren’t put in place, and there were no controls on who could update routing information, leading to a variety of issues with the database:

So everybody put their data in; they never removed old data, though. You’d look up a given prefix, and you’d see three different records, totally different information on who should be sourcing it, what the source AS is, who owns it, all different, and you just have to look at it and say, well that’s the most recent date stamp on that, so if they updated the date correctly – and originally in RADb the date was manually updated when you sent in the update, it wasn’t

²⁶See <http://www.ra.net/>, last retrieved Dec 11th 2013.

automatically added, that was eventually fixed – so a lot of those early dates were totally bogus too, because they forgot to update them. In fact, the one that was most recently updated was the one with the oldest date, sometimes! It was a mess; and it still is. [I17:4]

As the Internet grew, and ANS lost its dominant position, certain networks chose not to put their routing information into RADb, fearing that this would expose their private commercial information to their competitors:

The problem was that people didn't want to publish their policies, first off, because they thought that was revealing too much about the internal structure of their network, which they felt was proprietary at the time. And the other thing was, there was no real authentication, so anybody could put anything they wanted in it. So, like, [name redacted] registered 0/0. ... Back to the IRR thing, I think it was a good idea, and a lot of people still use IRR data, or try to use IRR data to generate filters. But the problem is, and always has been, that the data's stale, inaccurate, unauthenticated, all of the problems, right. [I10:5-6]

RADb and the IRR system continue to be used today, but not in the core role envisaged in the NSF call for the RA. As the topology of the Internet became more complex, and as economic rivalries emerged, it became ever more difficult to coordinate a centrally managed store of routing information.

The second strategy that the NSF adopted to ease the transition from the NSFNET to the Internet was to create a space for coordination amongst network administrators from both commercial and research networks, instituted as the North American Network Operators Group (NANOG). NANOG evolved from the “Regional-Techs” meetings of technical personnel representing the NSFNET's regional networks, who decided in 1994 to revise their charter to enlarge their representation.²⁷ In essence, NANOG acted as a mechanism to continue the culture of coordination and cooperation which was formed in the ARPANET, and carried on into the NSFNET. However, two key features differentiate NANOG from these ancestral Internet communities. First and foremost, the primary representation at NANOG was of operational personnel involved in the day-to-day running of their networks (as was Regional-Techs), not researchers, although NANOG remains open for anyone to participate. Second, economic interests began to play a strong role alongside operational interests, and these two interests sometimes came into conflict with one another, as the case of RADb discussed above illustrates. From an operational perspective, NANOG also became a key site for the production and reproduction of the social trust relationships needed to stabilize BGP.

Jon Postel continued to play the role of the IANA through the transition from the NSFNET to the Internet, although increasingly with staff to help him manage the IANA function. The IANA was eventually folded into the Internet Corporation for Assigned Names

²⁷From the NANOG history page, available at <http://nanog.org/history>, last retrieved Dec 11th 2013.

and Numbers (ICANN) in 1998. ICANN, along with five Regional Internet Registries (RIRs) - representing Asia, North America, Europe, Africa, and Latin America and the Caribbean - now manage the distribution of Internet numbers across the world. The constitution of the ICANN regime is a contentious issue, which I will return to in later chapters.

3.7 Institutionalizing the Internet

The analytical triad which I adopt for my analysis - standards development, resource allocation and topological form - became more clearly delineated as separate institutional forms as the NSFNET transitioned to become the Internet, while at the same time attempting to carry forward the spirit of the Internet - “We reject: kings, presidents and voting. We believe in: rough consensus and running code.” - in the functioning of these institutional forms. However, even as the functions for standards, resources and topology gradually separated and institutionalized independently over the course of the development of the Internet, they continued to influence one another.

These processes of separation and institutionalization as the Internet grew and became more complex significantly changed the nature of trust relationships amongst those involved in developing and operating the Internet. The early collegial social context of computer science researchers was a high trust environment, driven by thick relationships amongst researchers, and a tolerance for risk and uncertainty as new technologies were developed. This environment changed as the nascent Internet became larger, more complex, and more mission critical, with less tolerance for failure. The topology of the nascent Internet changed, from directly interconnected nodes on the ARPANET, to interconnected autonomous systems on the NSFNET, which grew to interconnect hundreds of research institutions, and eventually incorporated commercial entities as well. In order to manage the growth of the nascent Internet, the standards development and resource allocation functions also changed, formalizing and expanding the IETF and the IANA.

The interpersonal trust relationships of the early Internet were no longer sufficient to manage the ongoing growth and increasing complexity of the NSFNET, and then of the commercial Internet formed after the NSFNET was privatized in 1995. New organizational forms - such as the Regional-Techs group, and then NANOG - were created for coordination and collaboration amongst network administrators responsible for managing network interconnections, also functioning to develop trust relationships and practices of network administration. The functions of standards development and resource allocation were structured into the formal centralized governance institutions of the IETF and the ICANN regime. The norms of “rough consensus and running code”, these new organizational forms, and trust relationships together formed the basis for the distributed governance of Internet infrastructure. In the following chapters, I explore distributed governance in the modern Internet, by examining the evolution of these trust relationships and norms in centralized governance institutions, and in the practice of inter-domain routing.

Chapter 4

Understanding the Inter-Domain Routing System

There is little attention paid to the inter-domain routing system outside of the Internet's technical community. It is a component of the Internet's infrastructure which remains unseen, and largely unremarked upon in public discourse. In contrast, the Domain Name System (DNS) which maps human-readable names for machines (such as for websites) to IP addresses garners substantially more attention, if only because it represents the interface providing the user-friendly names that most people use to address machines on the Internet. If the DNS failed, we would lose the ability to address machines by name, effectively disabling access to websites and services across the Internet; but the Internet itself would continue to function normally. If the inter-domain routing system failed, the Internet itself would no longer exist, as the very interconnections between networks which constitute the Internet would cease to be.

The inter-domain routing system is perhaps the most critical component of the Internet's infrastructure, and it is essential to understand its functioning in order to unpack the opportunities and challenges to governance contained within it. In this chapter, I discuss the technology and economics of the inter-domain routing system in detail. I begin with a description of the Border Gateway Protocol (BGP), and the basic economics of inter-domain routing. I then go on to discuss the risks and uncertainties inherent in the technology and economics of the inter-domain routing system.

I analyze the risks and uncertainties in the inter-domain routing system in five broad categories. First, I look at the problem of prefix hijacking, which was how Pakistan took control of YouTube. Second, I put forward the problems of situated control and visibility, which limit network administrators' ability to respond directly to issues in other parts of the Internet which may affect the routing of their address space; again, the Pakistan/YouTube case is an example of this kind of problem. Third, I describe the hardware limits that networks have to address as the Internet grows and becomes more complex. Fourth, I engage with the problem of the mixed cooperation and competition which is inherent in the inter-domain routing system: networks must cooperate to interconnect, but at the same time

compete for customers. Finally, I address the problem of the exhaustion of IP address space, and what this means for the inter-domain routing system. These are the principal problems addressed by the system of distributed governance, which I engage with in the following chapters.

4.1 The Border Gateway Protocol

In order to understand the Internet's inter-domain routing system, the technology which enables it - BGP - must be examined in greater detail. BGP is one amongst a class of network protocols known as *routing protocols*. While most protocols in the TCP/IP suite are concerned with the transmission of data across a network, routing protocols are concerned with the discovery of appropriate routes for these data transmissions. As noted in the last chapter, these are broadly divided into interior gateway protocols for routing within an autonomous system, and exterior gateway protocols for routing between autonomous systems. Interior gateway protocols may be chosen and configured according to the needs of a particular autonomous system. In contrast, a common exterior gateway protocol must be used to interconnect a given set of autonomous systems. For the set of autonomous systems which make up the Internet, this common exterior gateway protocol is BGP.

In a packet-switched computer network, machines called routers act as the intermediate nodes which forward packets along the path from their origin to their destination. A router has multiple network interfaces which connect it to hosts and to other routers. Conceptually, these network interfaces are edges within the network, and the routers and hosts are nodes. Once a router receives a packet, its task is to decide which network interface should be used to forward the packet towards its destination.

Each router maintains a forwarding information base (FIB) for this purpose, indicating which network interface which should be used for a given IP address prefix. When the router receives an IP packet, it checks its FIB to determine the IP address prefix to which the destination IP address in the packet belongs. Since prefixes indicate a particular range of IP addresses, it is entirely possible for one prefix to be a subset of another. For instance, 192.168.1.0/24 is entirely contained within 192.168.0.0/16; in this case, the longer prefix (192.168.1.0/24, which has a 24-bit prefix) is said to be more specific than the shorter prefix (192.168.0.0/16, which has a 16-bit prefix). If a router finds multiple matching prefixes in its FIB for a packet, it will generally choose the most specific prefix to forward the packet. This is akin to a postal sorting center which has a box for letters to "Berkeley, CA", and another box for letters to "UC Berkeley, Berkeley, CA". Letters addressed to UC Berkeley will be put in the more specific box, rather than the catch-all "Berkeley, CA".

A router's FIB may be constructed through a variety of means, including manual configuration, information from a central repository, probing its neighbors for routes, or by learning routes announced by its neighbors. In general, the complete set of routes which a router acquires is much larger than the FIB, and is called the Routing Information Base (RIB).

A router's FIB is derived from its RIB through the application of algorithms and policies specific to the routing protocol in use.

Since there are over 480,000 IP prefixes currently visible in the inter-domain routing system, it is not practical for most routers to maintain a FIB containing all these prefixes.¹ The vast majority of routers maintain a FIB with a small set of routes for which some kind of routing policy is configured, alongside a special route to the address 0.0.0.0, called the *default route*. Packets which do not match any entry in the FIB are forwarded along the default route, on the assumption that the router at the other end of the default route will have better information in its FIB on how to forward the packet. Eventually, a packet must be forwarded through a router which has routes to the complete set of IP prefixes, without any default route. Such a router is called *default-free*, and the set of autonomous systems containing such routers are said to constitute the Internet's *default-free zone*. A routing table without a default route is called a "full table", or a "default-free table".

BGP routers function by announcing the prefixes that they can carry to their neighboring BGP routers, and learning routes from their neighbors through the same mechanism. BGP routers may also learn routes from the IGP within their autonomous system: since BGP manages routing between autonomous systems, packets destined for locations within an autonomous system will have to be forwarded to an IGP router. BGP uses the FIB to decide which routes should be announced to neighboring BGP routers. Routing policy set by administrators in BGP routers governs the application of rules to select routes from the RIB to the FIB, and to specify which routes from the FIB should be sent to which neighboring autonomous systems.

Routing policy is especially important for BGP routers, since it offers the only technological mechanism through which administrators at an autonomous system may exert local control over their immediate neighborhood. From a technical perspective, the inter-domain routing system is constructed and stabilized through the combination of local routing policy configurations. BGP routing policy provides certain explicit parameters which administrators may configure, alongside behaviors which are recommended by the BGP specification (Rekhter et al. 2006). The latter are the responsibility of BGP router vendors such as Juniper and Cisco, rather than of network administrators, although vendors will often provide the means for administrators to shape these behaviors to some degree.

BGP defines a limited set of messages which BGP routers may use to communicate with one another. These include messages to establish and maintain a BGP session between a set of routers, a message to report error conditions, and most importantly, a message to announce routes from one router to another. I shall be most concerned with this final kind of message - the BGP UPDATE message - since it provides the means through which inter-domain routing is enabled. At a minimum, a BGP UPDATE message must specify the following attributes:

- PREFIX: The IP address prefix being announced. A single UPDATE message may

¹For a report on the number of prefixes visible in the inter-domain routing system, see <http://www.cidr-report.org/as2.0/>, last retrieved January 8th, 2014.

contain multiple prefixes which share the same attributes.

- **ORIGIN:** Whether this prefix originated from within the autonomous system, from another autonomous system, or through some other means (such as a statically configured route).
- **AS_PATH:** The list of autonomous systems which this announcement has traversed. A BGP router which announces this route to BGP routers in neighboring autonomous systems is expected to prepend its autonomous system number to this attribute. This attribute is used for loop detection: if a BGP router discovers its own autonomous system number in the AS_PATH of an announcement, it knows that the announcement has already traversed its autonomous system, and therefore should be discarded.
- **NEXT_HOP:** The IP address to which to send traffic destined for the prefixes announced in this UPDATE message. This is typically the IP address of the router sending the UPDATE message.

BGP UPDATE messages may also include a variety of optional attributes, some of which are defined the BGP specification, and others which have been added as part of extensions to BGP. A BGP router is expected to simply ignore optional attributes which it does not recognize.

If a BGP router finds prefixes which are adjacent to one another it may be configured to aggregate them into a single prefix. For instance, the prefixes 192.168.0.0/17 and 192.168.128.0/17 may be aggregated into a single prefix 192.168.0.0/16. If prefixes are aggregated, the corresponding UPDATE message must include flags indicating that this route announcement was formed through aggregation. Conversely, administrators may de-aggregate an IP address prefix allocated to their autonomous system in order to specify different routing policy for different segments of that prefix, by announcing the de-aggregated prefixes in different UPDATE messages.

When a BGP router receives an UPDATE message, it must first determine the degree of preference for each prefix in the message, based on locally configured policy information. Multiple routes of equivalent degree of preference may point to the same destination, since multiple neighboring autonomous systems may be able to reach the same set of destinations. If so, the route with the shortest AS_PATH is preferred, on the assumption that a shorter AS_PATH represents a more optimal path through the inter-domain routing system.²

Once a BGP router finishes processing an UPDATE message, it must then relay any changes to its RIB to neighboring BGP routers. Different subsets of changes to its RIB are announced to different neighbors, depending on the relationships between autonomous systems. A simple categorization of relationships between autonomous systems separates them into peer and customer-provider relationships (see figure 4.1). A peer relationship typically involves no economic gain, on the assumption that peers exchange roughly equivalent

²Routes with equivalent AS_PATH lengths may be prioritized amongst using various optional attributes, which we need not concern ourselves with here.

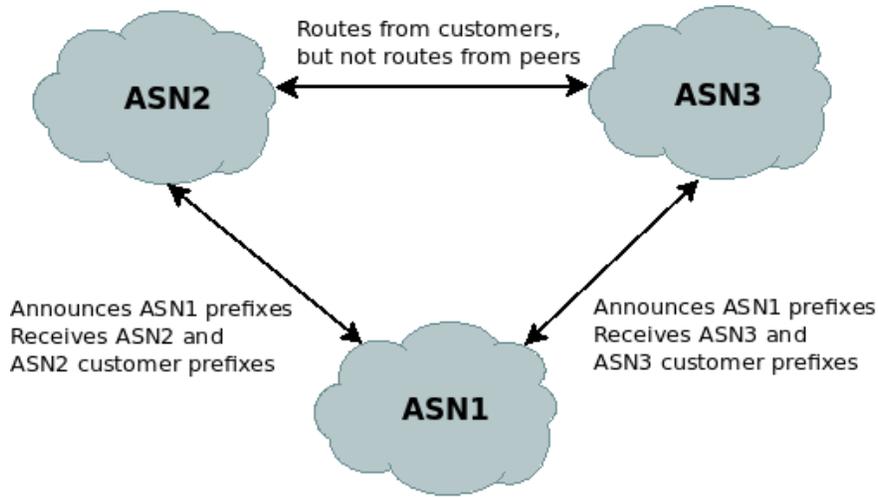


Figure 4.1: Peer-peer and customer-provider autonomous system relationships.

amounts of traffic. Customer-provider relationships involve customer autonomous systems paying their provider autonomous systems to carry their traffic. An autonomous system ASN1 which is a customer of provider autonomous systems ASN2 and ASN3 will announce its IP address prefixes to both these providers. However, ASN1 will be careful not to announce the routes it receives from ASN2 to ASN3, or those it receives from ASN3 to ASN2; if ASN1 were to do so, its network could potentially become an intermediary for traffic between ASN2 and ASN3. If ASN2 and ASN3 are peers, they will announce routes they receive from their customers to one another, but not routes from other peers. As with the ASN1 case, they will want to avoid becoming intermediaries for traffic which does not provide economic gain.

Economic relationships within the inter-domain routing system are more complex than this simple typology suggests, but this example serves to illustrate why autonomous systems filter the set of routes that they announce their neighbors. The result of these economic relationships is a topology which has a cartel-like arrangement of “Tier 1” autonomous systems that constitute the core of the Internet, together carrying default-free routing tables, typically with global infrastructure. Tier 1 autonomous systems are all “transit-free”, meaning that they do not pay any of their peers to carry their traffic; they are said to engage in “settlement-free” peering with other Tier 1 autonomous systems. Beneath these lie “Tier 2” autonomous systems which have regional infrastructure (e.g., within a country or within a state), contracting with Tier 1 providers to provide transit services to carry their traffic globally, but often engaging in settlement-free peering with other Tier 2 autonomous systems within the same region. Finally, “stub” autonomous systems occupy the edge of the Internet, establishing customer relationships with Tier 1 and Tier 2 providers to connect them to the larger Internet. These arrangements have been complicated in recent years with the rise of content distribution networks (CDNs) which attempt to locate content closer to cus-

tomers, and so form an overlay with multiple geographically distributed points of presence, and interconnections with autonomous systems at all levels.

Each autonomous system has a different routing table, with these differences conditioned both by the size of the autonomous system, and by its position within the Internet. A Tier 2 or stub autonomous system will carry only a partial routing table with some portion of global IP address space routed via default routes carried by an upstream Tier 1 network which will maintain a complete default-free routing table. In addition, each autonomous system has a routing table formed based on its own situated perspective within the Internet, since routing tables are formed from routes that are visible to a specific autonomous system via its neighbors, rather than from some common global map of routes. As routing updates propagate through the inter-domain routing system, individual autonomous systems' routing tables will converge upon a common representation. However, ongoing updates to routing information ensure that convergence is never complete, but rather a continuous process. The *convergence time* required for the results of a given routing update to propagate across the Internet is a critical measure of the efficacy of the inter-domain routing system, and indeed, of any routing protocol. A longer convergence time implies a greater period over which routing tables in different autonomous systems will be inconsistent with one another, leading to potential routing faults.

4.2 Risk and Uncertainty in the Inter-Domain Routing System

In this section, I explore the risks and uncertainties inherent in the Internet's inter-domain routing system, which are the primary objects of coordination and collaboration amongst the Internet's technical communities. These are wide-ranging, spanning across technical features of BGP, issues in the topological arrangements of networks, and the economics of the inter-domain routing system.

4.2.1 Routing by Rumor

As I noted in the last chapter, BGP provides no mechanisms for evaluating the veracity of claims in routing announcements. I argued that BGP took this form for two reasons. First, the extremely trusting and tight-knit culture of the research community operating the NSFNET which was not immediately concerned with security. Second, the hierarchical topology of the NSFNET which provided for centralized points of control, which acted as an easy check on spurious routing announcements. The Internet, however, exposes no such centralized points of control at a global level.³ Given the problems inherent in reliably evaluating BGP announcements, the process of routing in BGP is sometimes facetiously

³Although such centralized points of control may be implemented via carefully constructed local network topologies.

referred to as “routing by rumor”.⁴ In addition,, the arrangement of interconnections amongst networks on the Internet must be understood in terms of contracts and broader economic and political arrangements, as much as in terms of the nature of the technical communities involved in day-to-day network operations. As one of my interviewees commented:

And in the early days, I think, that was probably modeled after the NSFNET backbone where there was a group of people operating the network, and they were all the same, in some sense were all in the same group. So, this thing that developed, where basically people running the network don’t even know each other, isn’t something they probably envisioned. [I10:4]

An autonomous system may trust the routing updates from its immediate neighbors for IP address space which they are allocated, given a contractual relationship for interconnection. Routing announcements relayed from neighbors of these neighboring autonomous systems are much more difficult to evaluate. An administrator at a campus network related these difficulties to me:

If I have a direct peering with you, you and I can make up some pretty strict rules about what we’ll pass back and forth. If I’m dependent on you to reach the world, I pretty much have to say, give me everything you’ve got, and I have to accept it, whether I trust it or not, and that’s really the model that we have today. ... But, once you get close to the default-free Internet, you essentially have to accept whatever announcements someone sends you. The only ones you can really filter are the ones you get directly from them, and of course, the ones you send back. Other than that, you have to trust that they’re sending the correct information. [I1:10]

A campus network, topologically positioned as a stub, is entirely reliant on its upstream network providers for routing information. It may maintain filters at its BGP routers to ensure that it will only announce its own IP address space, but it is incapable of applying similar filters to validate routing information from its upstream network providers. These issues are magnified manifold amongst autonomous systems which act as transit providers, often carrying traffic originating and terminating at autonomous systems with which they do not have a direct contractual relationship. An administrator at a Tier 1 network expanded on these concerns:

Well, the first thing is, if you’re looking around the network, you see a prefix being originated from a particular ASN. How do you know that’s really legitimate? Maybe the person who’s at the border wants to multi-home, maybe they have

⁴One of my interviewees pointed me to his use of the phrase [I24:15] in a widely referenced IAB statement from 2010, available at <http://www.ietf.org/mail-archive/web/ietf-announce/current/msg07028.html>, last retrieved Jan 24, 2014. This phrase may be found in a variety of locations, including in textbooks from Cisco Systems, a leading networking equipment vendor, dating as far back as 2001: see <http://www.ciscopress.com/articles/article.asp?p=24090&seqNum=3>, last retrieved Jan 24, 2014.

two different transit providers, and each one likes to advertise differently. It gets really hard to know what is legitimate. If you tie it into the AS path, people could spoof an AS, they could put a different AS on there. It's really hard to say what is real and what is not. From a customer, it's another story, right, with a customer there's the assumption that you've already done your homework, you've already done the due diligence, saying this customer's authorized to advertise this prefix, but when it's through a peer, that's a little bit harder. [I18:7]

An autonomous system may announce address space which is not allocated to it. If its upstream network providers do not filter its routing announcements correctly, these spurious announcements may leak to the wider Internet and enter the default-free zone. This may happen for exactly the same IP address prefix as another autonomous system, but it may also be for a more specific portion of that prefix. In the former case, both routes are in contention with one another, with one or the other preferred based on local policies at intermediate autonomous systems relaying these announcements. In the latter case, the more specific prefix is preferred by default, and will almost certainly result in the redirection of data traffic. Such an event is known as a “prefix hijack” or a “route leak”, and is exactly what happened when the Pakistan Internet Exchange hijacked YouTube's address space. These are not unusual occurrences, and may happen either through malicious intent, or through inadvertent misconfiguration, or software faults. Malicious prefix hijacking events are driven by the intent to masquerade as a particular IP address prefix, thereby capturing all traffic directed to it, and being recognized as the legitimate source of traffic from that IP address prefix.⁵ Prefix hijacking can be difficult to detect and quantify: one study detected between 26 and 95 possible prefix hijacks in a single month, with many more possible misconfigurations (Boothe et al. 2006), while another study suggests that there have been between 5 and 20 prefix hijacks every year between 2003 and 2010 (Khare et al. 2012).

A famous instance of large-scale prefix hijacking which has gone down in Internet lore is the case of AS7007. On April 25th, 1997, routes to all IP address prefixes on the Internet briefly pointed to AS7007, in effect disabling the inter-domain routing system. The AS7007 border routers received a default-free routing table of about 23,000 routes from one of its customers. Due to an apparent software fault, and possible misconfiguration, the AS7007 border router then deaggregated the IP address prefixes in these routes - increasing the number of routes to over 73,000 - stripped out the AS_PATH, and then relayed these announcements to all of its peers. Even after rebooting its border routers, and eventually disconnecting itself from the Internet, these routes continued to be viewed as authoritative,

⁵More sophisticated variants of prefix hijacking may function as “man-in-the-middle” attacks, in which traffic - once captured - is relayed to the legitimate holder of the hijacked IP address prefix. For an example, see <http://www.renesys.com/2013/11/mitm-internet-hijacking/>, last retrieved Aug 6, 2014. Prefix hijacking may also be used to defeat email spam filters by intermittently announcing and withdrawing IP address prefixes, sending email from this intermittent address range only in the period during which it is announced (Ramachandran and Feamster 2006).

likely because of intermediate autonomous systems which kept relaying them to their peers and customers. The person responsible for AS7007, Vincent Bono, eventually took to the NANOG email list to explain the event as best as he could, and apologize to the NANOG community at large who were all affected by this incident.⁶ Such an event is called a “full table leak”. While these are far less common than prefix hijacks, they do very occasionally still occur, albeit at a regional - rather than global - scale.⁷

Prefix hijacking is only part of the story, however. The autonomous system originating a routing announcement, or any intermediate autonomous system relaying the announcement, may be “multi-homed”, announcing the route to multiple upstream network providers. Intermediate autonomous systems may therefore receive multiple routing announcements claiming to be able to carry traffic to the same IP address prefix and origin autonomous system through different paths, as represented in the routing announcement’s AS_PATH attribute.

Autonomous systems may manipulate the AS_PATH in routing announcements. A commonplace manipulation is to use “AS prepending” to insert the same autonomous system number multiple times, causing upstream networks to prefer the route in that announcement less, on the assumption that it represents a longer route to the originating autonomous system. This is commonly used for traffic engineering, to signal route preferences to upstream networks which are potentially more than one degree of separation removed. This kind of manipulation may be carried out by the originating autonomous system, or any intermediate autonomous system involved in relaying the routing announcement.

Just as with prefix hijacking, this kind of manipulation may also be employed to hijack traffic. An autonomous system may craft an announcement purporting to be from the legitimate assignee of an IP address prefix, by including the legitimate assignee’s autonomous system number at the end the AS_PATH. There is no easy way for any other autonomous system receiving this announcement to evaluate whether or not the legitimate assignee of the IP address prefix announce is actually a customer of the autonomous system issuing this routing announcement. The situation becomes even more complex when we consider that spurious autonomous system numbers may be embedded at any position within the AS_PATH. These are real concerns: in 2013, networks in Iceland and Belarus were observed to be redirecting traffic to themselves before handing it onwards to the destination for which it was intended.⁸ A respected member of the NANOG community claims to have observed over

⁶See <http://www.merit.edu/mail.archives/nanog/1997-04/msg00444.html>, last retrieved Jan 22nd, 2014.

⁷For information on a nearly-full table leak confined to Brazil in 2008, see <http://www.renesys.com/2008/11/brazil-leak-if-a-tree-falls-in/>, last retrieved Jan 22nd, 2014. For information on a full table leak from Turkey in 2004 which did affect the global Internet, see <http://www.renesys.com/2005/12/internetwide-nearcatastrophela/>, last retrieved Jan 22nd 2014. Renesys is an Internet market research firm which actively monitors the inter-domain routing system from over 400 autonomous systems around the Internet, and is actively involved in disseminating portions of their research publicly to the Internet’s technical communities.

⁸See <http://www.renesys.com/2013/11/mitm-internet-hijacking/>, last retrieved Jan 22nd, 2014.

11 million such events since 2008 through an analysis of publicly available routing data.⁹ In order to accurately evaluate the veracity of a routing announcement, an autonomous system must be able to validate both that the originating autonomous system is the legitimate assignee of the IP address prefix being announced, and that all intermediate autonomous systems relaying the announcement (visible in the AS_PATH) are actually interconnected with one another.

The Routing Assets Database (RADb) and the Internet Route Registries (IRRs) described in the last chapter were meant to stave off some of these issues. However, their capabilities are limited, lacking participation by all autonomous systems on the Internet. More importantly, the entities operating these databases are incapable of enforcing their use, since centralized topological positions do not exist within the inter-domain routing system, from which this kind of enforcement can proceed on a global scale. A network administrator from a Tier 1 autonomous system related his concerns with RADb and the IRRs to me for network operations in the North American region:

There's a number of problems with it. The data is not well-maintained. It's partly a function of the fact that people stuff data in it to make sure routing doesn't break, and then they never update it at any point in the future, sort of abandoned. It's also problematic that anyone who is a registered user can make updates to any route. I mean, I can put stuff in on behalf of my customer, I can put stuff in on behalf of my customer for their other upstream ISP, so there's not a lot of integrity there. And the people who run the database have no relationship with the people who hold the addresses. [I14:6]

The maintenance of accurate data within RADb and other IRRs continues to be a problem. For instance, on a recent thread on the NANOG email list,¹⁰ a network administrator comments that IP address prefixes allocated to his autonomous system are registered by the ISP Cogent in RADb, and to a number of other ISPs in the IRR maintained by the ISP Level 3. The responses suggest that it is fairly straightforward to get the entries in RADb changed, but that Level 3 “fills their IRRDB up with piles of crap”.¹¹

There also concerns that IRR information may be used by competitors to gain information about an ISP's customers, and that the maintenance of IRR entries represents an unreasonable administrative overhead, as noted in RFC 4277:

The efforts of the IRR participants has been severely hampered by providers unwilling to keep information in the IRR up to date. The larger of these providers have been vocal, claiming that the database entry, simple as it may be, is an

⁹See <http://mailman.nanog.org/pipermail/nanog/2013-December/062607.html>, last retrieved Jan 24, 2014.

¹⁰See <http://mailman.nanog.org/pipermail/nanog/2014-January/063550.html>, last retrieved Jan 24, 2014

¹¹See <http://mailman.nanog.org/pipermail/nanog/2014-January/063551.html>, last retrieved Jan 24, 2014.

administrative burden, and some acknowledge that doing so provides an advantage to competitors that use the IRR. The result has been an erosion of the usefulness of the IRR and an increase in vulnerability of the Internet to routing based attacks or accidental injection of faulty routing information. (McPherson and Patel 2006:15)

In consequence, RADb and the IRRs are used only by a limited set of the autonomous systems on the Internet. However, it is important to note that the IRRs maintained by Regional Internet Registries (RIRs) are considerably more reliable than RADb and IRRs maintained by individual ISPs, quite simply because the RIRs are responsible for the allocations which are represented in their IRRs. I will return to this issue later, when I discuss ICANN regime for the allocation of IP address space and other critical Internet resources.

4.2.2 Situated Visibility, Situated Control

The inter-domain routing system is premised upon the assumption that network administrators can act as they wish within the bounds of their own autonomous systems. It follows that the authority of network administrators to act as they wish is also *limited* to the bounds of their own autonomous systems. When faced with issues at a neighboring autonomous system, an administrator may have a clear line of accountability to invoke, since the relationship with the neighbor is likely warranted through a contractual arrangement. However, when faced with issues at autonomous systems further removed within the inter-domain routing system, it can be more difficult to have the necessary remedial action taken.

These issues could involve prefix hijacking or AS_PATH manipulations. They may also be related to degradation or loss of a particular network path. Such issues can be difficult to diagnose, since an administrator's perspective on the inter-domain routing system is always *situated*, with visibility originating from, and conditioned by, the relative position of his or her autonomous system within the inter-domain routing system. This formed a theme underlying many of my interviews, and arises regularly on the NANOG email list. I will illustrate these kinds of issues with recent examples drawn from the NANOG email list in December 2013 alone. It is worth noting again here that the NANOG email list is one of the most highly visible email lists in the global network administration community, with publicly available archives. When someone sends an email to the list, they do so with the knowledge that they are entering into conversation with a global community.

Crocker Communications, which uses Cogent Communications and Sprint for their upstream connectivity, recently had an issue connecting to Level 3 through Cogent, although Level 3 was still reachable through Sprint. They sent email to the NANOG mail list asking whether anyone else was having similar issues.¹² While Cogent initially claimed that this was the result of a capacity issue with Level 3, it later emerged that the issue arose because Co-

¹²See <http://mailman.nanog.org/pipermail/nanog/2013-December/062638.html>, last retrieved Jan 24, 2014.

gent had incorrectly applied a filter to their interface to Crocker Communications.¹³ While the problem could have been with Level 3's connection with Cogent, this example shows how it can be difficult to resolve issues even with immediately neighboring autonomous systems.

Administrators may ask whether others are experiencing the same issues that they are, to establish whether the problems they're observing are related to their immediate circumstances, or are representative of a broader problem. For instance, an administrator noted issues with network capacity provided by Abovenet on links between Delaware and New Jersey, and between Delaware and California,¹⁴ proceeding to ask whether others were also experiencing the same problems.

Others may ask for contacts at specific autonomous systems to help resolve issues that they're observing. Consider the administrator who noticed a degradation of connectivity at sites connected through Level 3, and asked for contacts at Level 3 or Comcast, proceeding on the hypothesis that the cause of the problem was network congestion between Level 3 and Comcast.¹⁵ A similar message noted issues with network reachability to the website at www.army.mil, which is hosted on AS1503, and asked for a contact at AS1503 to help resolve the problem.¹⁶

I have restricted my examples to issues related to connectivity and inter-domain routing. However, the same kinds of requests are visible on the NANOG email list for a variety of other issues, including email services and the Domain Name System (DNS).

The issue of the visibility and nature of relationships amongst autonomous systems has also been an ongoing concern for researchers (Gao 2001; Li et al. 2004; Oliveira et al. 2008). Even with observations of routing announcements from distributed locations around the Internet, it is difficult to discover the complete graph of relationships amongst all autonomous systems. Some relationships can at best be inferred, given that they may not be visible from the vantage points at which routing announcements are collected; between 10% and 85% of all relationships may not be visible from routing announcements (Oliveira and Willinger 2010). Researchers find this problem of great importance, since accurate maps of the inter-domain routing system are needed in order to be able to model the behavior of new Internet protocols. This is also an important problem for network operators, who would find such maps useful in determining the location and cause of routing faults.

¹³See <http://mailman.nanog.org/pipermail/nanog/2013-December/062719.html>, last retrieved Jan 24, 2014.

¹⁴See <http://mailman.nanog.org/pipermail/nanog/2013-December/062907.html>, last retrieved Jan 24, 2014.

¹⁵See <http://mailman.nanog.org/pipermail/nanog/2013-December/063283.html>, last retrieved Jan 24, 2014.

¹⁶See <http://mailman.nanog.org/pipermail/nanog/2013-December/062726.html>, last retrieved Jan 24, 2014.

4.2.3 Hardware Limits

The capacity of the inter-domain routing system in individual autonomous systems is constrained by the capabilities of the BGP router hardware and software in use. There are two significant hardware limits involved in establishing these capabilities: the amount and speed of memory, and processing speed. The amount of memory a router requires is determined by the size of its FIB, RIB and filtered routing tables that it announces to neighboring routers. The speed of the memory is primarily determined by the number of packets that the router handles per second. A higher throughput rate for packets implies a need for faster memory, since the router must perform a lookup for each packet on the copy of the FIB in memory in order to decide the network interface over which to forward the packet. Processing speed for a router depends on processing requirements for the router's two primary components: the data plane which is responsible for forwarding packets, and the control plane which is responsible for managing the routing tables. As with memory, a higher throughput of packets per second implies a need for faster processing for the data plane. A higher number of routing announcements, and greater complexity of routing policies, imply a need for faster processing at the control plane.

The “big iron” routers capable of handling high volumes of data traffic and carrying a default-free RIB are extremely expensive. This is in part due to the progression from routers using generic processors with specialized routing software, to routers which embed routing logic into hardware using ASICs (Application Specific Integrated Circuits) to provide greater processing speeds. The high speed memory chips required to provide fast access to the RIB are also major contributors to the cost of routers. The number of prefixes being announced to the default-free zone has increased steadily over the years (see figure 4.2). This is an ongoing cause for concern, as increases in RIB size also imply an expensive corresponding need for faster processors and faster memory in BGP routers. This issue is a common subject of discussion amongst the Internet's operational communities, since this implies a need for capital investment which is often outside the immediate purview of the network administrators responsible for managing BGP routers.

Many autonomous systems deaggregate their allocated IP address prefixes into sub-prefixes which they assign to customers, or use to direct their traffic along different routes. The more prefixes are deaggregated, the larger the size of the default-free routing table, since each deaggregated prefix is announced separately. Autonomous systems should aggregate their prefixes to the extent possible, but this does not happen in practice. If the 489,912 prefixes currently visible in the default-free routing table were completely aggregated to the extent possible, the size of the table would reduce to 275,505 prefixes.¹⁷ This is primarily a problem for autonomous systems in the default-free zone, although other autonomous systems may be affected as well in the instance of a full table leak, such as those in the incidents detailed earlier. One approach to addressing these issues is to limit the number of routes that a router may accept. However, this approach may in itself result in a faulty routing table, since there is no way to distinguish between good and bad routes when deciding which

¹⁷From <http://www.cidr-report.org/as2.0/>, last retrieved Jan 23, 2014.

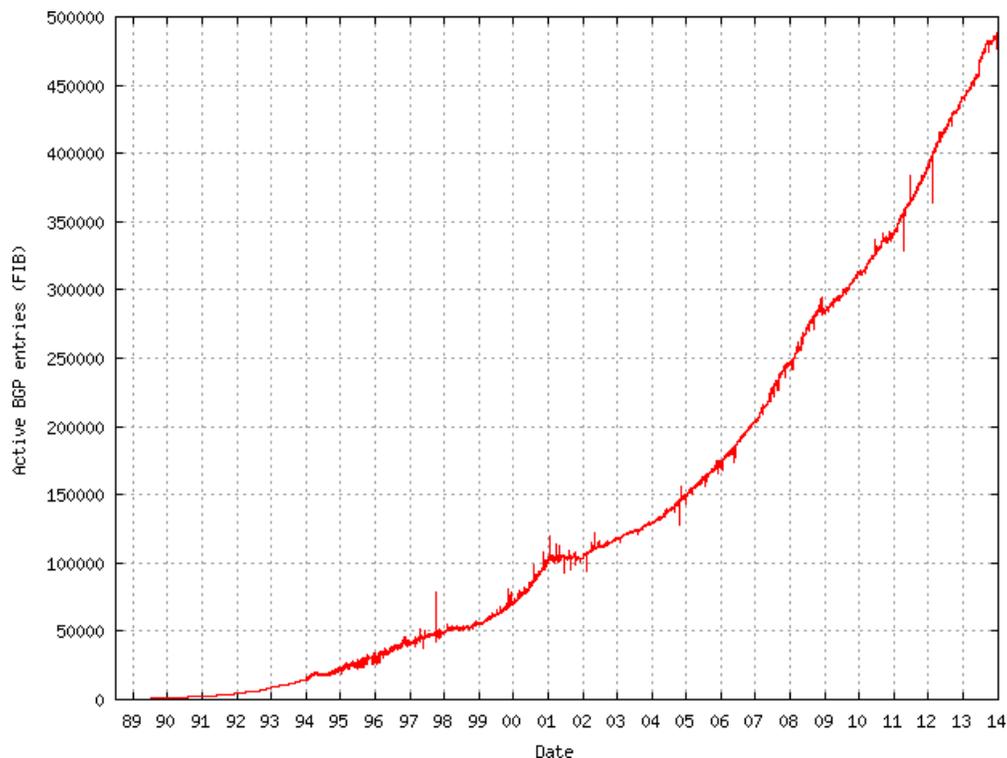


Figure 4.2: The size of the Internet’s default-free zone over time, from the CIDR Report, available at <http://www.cidr-report.org/as2.0/>, last retrieved Jan 23, 2014.

routes to drop from routing announcements for a full table. These kinds of issues place a load on both the data plane and the control plane of BGP routers, since an increased number of routing announcements need to be processed, and lookups on the FIB become more time consuming as the FIB grows in size. As a network administrator at a campus network put it to me:

We’re far enough from the core... I don’t think we’ve ever been slammed with a full table leak. There’s always been someone upstream of us who’s been slammed with the full table leak, and they took care of the issue for us. That is one of the things about being a campus, we’re far enough from the core, you’re far enough from the core that that kind of stuff usually doesn’t reach you. We’ve seen a couple of cases, and actually not that long ago, probably within the last 4 years, where somebody has announced a horrifically large number of routes, and typically it will hit somebody else’s memory limits before it gets down to us, and hits ours. So we don’t get the immediate effect, we see it the next day in the configured routes and we say, oh look at that, that was a lot of routes. There are a lot of sites that have protection in place, and they say, oh, you know, I am getting 100,000 routes today, I will have a limit of 150,000. Well, that may

have been fine 4 or 5 years ago, but it's not today, we get over 250,000 routes today, and it just keeps going up, it makes no sense to have those kinds of limits. Besides, you may protect your router from falling over in that case, but basically, if the number of routes goes up that high, you're breaking connectivity, because there's nothing in the router that says, drop this route and not that route. So when you drop routes, you dropping good routes as well as bogus routes. You've basically broken the service that you're trying to run anyway... [I1:11]

BGP routers may also be faced with routes from neighbors which oscillate between being announced, and then being withdrawn. This phenomenon is called “route flap”, and can impose severe penalties on a router's control plane as it has to continually process announcements for the same route. Route flap can have effects beyond immediately affected routers, since they in turn will announce these routes to their neighbors, and these neighbors to their neighbors, and so on. As a result, route flaps can affect all autonomous systems in the default-free zone, and may potentially affect the entire inter-domain routing system. Apart from the additional processing load on a router's control plane, route flaps also increase the convergence time for the default-free routing table, resulting in increased inconsistency of routing information amongst autonomous systems. This has been a well-known issue since the transition from the NSFNET to the Internet, with the IETF issuing RFC 2439 in 1998 detailing methods for damping route flap (Villamizar et al. 1998). Route flap remains an ongoing concern for the inter-domain routing system, with some autonomous systems issuing thousands of updates per day for their prefixes.¹⁸ Since BGP is a policy-based routing protocol which allows autonomous systems to construct locally optimum conditions, rather than configuring to a global optimum, it is even possible that routing tables in the inter-domain routing system will *never* converge (Feamster et al. 2004) under certain topological arrangements. Researchers have attempted to develop guidelines for local BGP configurations which would guarantee route convergence (Gao and Rexford 2001), but it is far from clear how well individual autonomous systems might conform to such requirements.

Attacks on a BGP router's data plane may also affect its control plane, since these typically share the same physical medium (Zhang et al. 2007). When targeted carefully, such attacks have the potential to affect large swathes of the inter-domain routing system (Schuchard et al. 2010).

4.2.4 Economic Organization

The inter-domain routing system is characterized by incentives which are at odds with one another. Autonomous systems must cooperate to interconnect, yet at the same time they must compete for customers and better terms of interconnection. It is only through interconnection that autonomous systems can offer value to their customers, by providing them access to all possible destinations on the Internet via other autonomous systems. Yet the

¹⁸See <http://bgpupdates.potaroo.net/instability/bgpupd.html> for current data, last retrieved Jan 23, 2014.

terms of interconnection themselves often come under dispute, as the corporations operating autonomous systems compete with one another to maximize profits. Piscitello and Chapin (1993:421-422) put it well:

Inter-domain routing ... plays the paradoxical role of facilitating communication amongst open systems for which communication is a (politically) sensitive activity ... that can produce highly counter-intuitive answers to what look like simple technical questions.

They go on to suggest, tongue-in-cheek, that “the only large scale inter-domain routing protocol that is likely to be deployed in the near future will be implemented as an army of lawyers on bicycles” (Piscitello and Chapin 1993:423), making reference to the complexity of peering agreements, the technical/commercial contracts governing the interconnection of networks.

The stub autonomous systems at the peripheries of the Internet are rarely part of this dynamic, since they simply buy transit from one or more upstream autonomous systems which provide them access to the broader Internet. In conversations with network administrators at smaller ISPs about their BGP configurations, I have often been told that they only have one upstream network provider, and rarely need to think strategically about managing their upstream connectivity. However, even stub autonomous systems may engage in strategic behavior to reduce their transit provider costs. For instance, UC Berkeley is part of a coalition of educational institutions across California which obtain connectivity to the Internet through the Corporation for Education Network Initiatives in California (CENIC),¹⁹ which operates the California Research and Education Network (CalREN). UC Berkeley has direct connectivity to other educational institutions in California through CalREN, saving considerably on transit costs. In addition, by aggregating demand, CENIC acts as a Tier 2 autonomous system which can negotiate favorable interconnection agreements with other Tier 2 and Tier 1 autonomous systems.²⁰

The competition/cooperation dynamic plays out more aggressively amongst the Tier 1 and Tier 2 autonomous systems involved in providing transit to one another, and to stub autonomous systems. Tier 1 autonomous systems form a relatively closed club, to which most privately-run Tier 2 autonomous systems aspire. To be a Tier 1 autonomous system, after all, is to eliminate the need to buy transit from other autonomous systems, since Tier 1 autonomous systems all engage in settlement-free peering with one another. However, the requirements to achieve status necessary to engage in settlement-free peering with Tier 1 or Tier 2 autonomous systems can be onerous and expensive.

All autonomous systems which offer transit services have well-defined conditions under which they will accept another autonomous system as a settlement-free peer. Verizon Busi-

¹⁹See <http://www.cenic.org/>, last retrieved Jan 27, 2014.

²⁰The classification of autonomous systems as Tier 1, Tier 2, and stub is something of an oversimplification. However, I will continue to use these terms, as they present a useful set of categories for the purposes of the discussion here.

ness' policy in this regard are typical of most other Tier 1 autonomous systems.²¹ This policy specifies that a potential settlement-free peer must be able terminate traffic in at least 50% of the geography covered by Verizon Business' networks, that the ratio of volumes of traffic exchanged must no exceed 1.8:1, that certain minimum amounts of network capacity must be available for interconnection, that certain minimum amounts of autonomous systems should be transit customers, and that the aggregate amount of traffic should exceed a specific minimum level, in addition to various operational requirements. As this list of requirements makes clear, any potential settlement-free peer must have reached a critical size in terms of customers, geographic reach and traffic volumes, which is in itself a capital-intensive process. Once a settlement-free peering agreement has been arrived at, it is by no means cast in stone; I will describe later how high bandwidth services such as video streaming can be the cause of challenges to these agreements.

In some cases, it be quite literally impossible to enter into a settlement-free peering relationship with a particular autonomous system. Several of my conversations dealt with how Tier 1 autonomous systems have a very simple and short peering policy: "No!" [F-NANOG58:4-5,136, I17:22]. Any autonomous system offering transit services would, after all, prefer to have a paying customer rather than a settlement-free peer. Depending on the goals of the entity operating an autonomous system, it may often be simpler and cheaper to buy transit, rather than build out the capacity to become a settlement-free peer. One of my interviewees commented on these often contentious relationships:

I think you basically get to this situation where - which people hate - but the large providers all have the capability of destroying each other's networks by mistakes by bad management. ... If you do bad capacity planning, if you make mistakes, then your peers will get cranky. If you're a transit-free network, then what are people gonna do, de-peer you? But if you're anywhere smaller than that your whole reputation then, depends upon your not pissing off your peers. [I23:12]

In the tiered model of the Internet, capital flows from the autonomous systems at the peripheries of the Internet towards the core. A stub autonomous system such as UC Berkeley pays CENIC, which in turn pays various Tier 1 providers to carry traffic. UC Berkeley has a mix of traffic, since it provides Internet access to students, staff and faculty, and also hosts the various web servers, mail servers, and so on, which establish UC Berkeley as a presence on the Internet. Other autonomous systems which provide commodity Internet service to ordinary users - small and medium ISPs such as Sonic.net - similarly pay upstream network providers for transit services, but act primarily as a destination for data, as their users request content of various types via their web browsers, email clients and so on. Some Tier 1 autonomous systems - such as AT&T, Comcast and Verizon - also provide commodity Internet service in addition to transit services. In general, capital flows from the peripheries of the Internet - stub autonomous systems and commodity Internet users - towards the core.

²¹See <http://www.verizonenterprise.com/terms/peering/>, last retrieved Jan 27, 2014.

The process of growing from Tier 2 to Tier 1 status, or even of maintaining Tier 1 status, is often fraught. Autonomous systems on both sides of a settlement-free peering connection monitor each other's capacity and behavior, with ongoing negotiations which can sometimes result in peering conflicts: incidents in which autonomous systems disconnect, or "de-peer", from one another. This is often part of a routine process of expansion, as an autonomous system determines that it has grown sufficiently that its peers no longer fulfill settlement-free peering criteria, and may be converted to transit customers. However, when such incidents occur amongst autonomous systems which rely on each other to get to particular destinations in the inter-domain routing systems, this can result in a partition of the Internet, a condition in which certain IP address prefixes are unreachable from one another. This is a condition which results in the Internet being split into two or more components with limited connectivity to one another: multiple internets, rather than a single global Internet. As extreme as this sounds, it has nonetheless actually occurred in recent years.

An ISP called Cogent Communications has been at the center of several peering disputes over the last few years, as it has tried to grow towards Tier 1 status. Cogent has been involved in peering disputes with the American ISP Level 3 in 2005,²² and with American ISP Sprint,²³ and the European ISP Telia in 2008.²⁴ In each of these cases, the result was a partition of the Internet, as customers who were single-homed behind Cogent were no longer able to connect to customers single-homed behind Level 3, Telia or Sprint, and vice versa. Although all of these incidents were eventually resolved, through negotiation and customer pressure, they did last for multiple days. Unlike technical failures in inter-domain routing which are measured in minutes or hours, economic disputes occur over a sufficient period of time that they become noticeable to customers, and even the mainstream press. Peering disputes can be characterized as the market simply working itself out, and for the most part, this is true. However, when peering disputes lead to de-peering which partitions segments of the inter-domain routing system, this is perhaps better understood as the market failing to preserve the stability of the global Internet.

As the types of content available on the Internet have become more bandwidth-intensive - such as audio and video streaming - a new category of autonomous system has emerged within the inter-domain routing system: the Content Delivery Network (CDN). CDNs function by moving content as close to those requesting it as possible, to reduce the costs of distributing this content to their customers. These may be specialized CDN companies such as Limelight or Akamai which cache content for their customers at locations around the world. These may also be Internet services companies such as Google which may cache content (such as YouTube video) closer to those viewing it (Gill et al. 2008). The result is

²²See <http://www.networkworld.com/edge/news/2005/102805-cogent-level3.html>, last retrieved Jan 27, 2014.

²³See <http://www.renesys.com/2008/10/wrestling-with-the-zombie-spri/>, last retrieved Jan 27, 2014.

²⁴See <http://www.renesys.com/2008/03/you-cant-get-there-from-here-1/>, last retrieved Jan 27, 2014.

a new topological feature, an autonomous system which has global reach (for geographically distributed content caches), but does not offer transit services. CDN services save their operators and customers money since they reduce the amount of bandwidth needed by their customers in order to deliver their content. For example, lacking a CDN, Google would have to budget for higher bandwidth costs, as it would have to pay its upstream network providers to carry every request for a video to its destination. Bandwidth-intensive services often employ a combination of paying for transit and CDN services to deliver their content.

CDNs and bandwidth-intensive services pose a challenge to the tiered model of the Internet. As CDNs push the locations of content towards the periphery of the Internet, autonomous systems providing transit services see a reduction in bandwidth requirements (and corresponding revenue) from their customers, as customers may get particular kinds of data directly from a CDN, rather than needing to have it flow through their transit provider (Labovitz et al. 2010). In addition, autonomous systems which are the primary transit providers for bandwidth-intensive services may find themselves in conflict with autonomous systems which control the commodity Internet users (the “eyeballs”) requesting these bandwidth-intensive services.

In November 2010, Level 3 and Comcast engaged in a peering dispute as a result of this kind of conflict. Level 3 had just become the primary transit provider for Netflix, resulting in a traffic imbalance between Level 3 and Comcast, as users of Comcast’s commodity Internet service requested bandwidth-intensive video streams from Netflix. This was complicated by the fact that Comcast also offers video content via their cable service: competing with Netflix, while at the same time delivering Netflix to their customers. Comcast claimed that they could no longer engage in settlement-free peering with Level 3 due to this traffic imbalance. Level 3 responded by suggesting that net bandwidth exchanged is no longer a good measure for evaluating the status of peers. Instead, they suggested a new measure, “bit-miles”, which calculate the distance that data is carried by an autonomous system, rather than simply the amount of traffic exchanged. Level 3 could find an alternate path to Comcast through Tata Telecommunications; however, Comcast ran their ports with Tata at capacity, resulting in significantly degraded traffic through this path. In the balance, Comcast controlled access to their commodity Internet customers, while Level 3 needed to reach these customers in order to provide the service that Netflix paid it for. As one commentator put it, “content is mobile, but eyeballs are not.”²⁵ Content may take a different path, but the eyeballs looking at this content are captive behind their Internet service provider. Comcast and Level 3 eventually reached an accommodation, although the terms of this accommodation are sealed under a Non-Disclosure Agreement (NDA), which is a common practice for peering agreements. A report on the Comcast-Level 3 dispute from the US Federal Communications Commission (FCC) notes that transparency is a critical issue here: regulatory agencies need to be able to force disclosure of the terms of peering agreements, in order to perform antitrust evaluations of these new business models (Rose 2011).

²⁵See <http://www.internap.com/2010/12/02/peering-disputes-comcast-level-3-and-you/>, last retrieved Jan 28, 2014.

Regardless of the arguments on both sides, this conflict raised the issue of “network neutrality”. Simply put, this is the notion that autonomous systems should not prioritize traffic by type or source; that criteria for evaluating interconnection agreements should not penalize new kinds of traffic (such as video and audio streaming, or IP telephony), or new entities offering Internet services. In both cases, the flow of capital on the Internet would change from a periphery-to-core model to one in which autonomous systems which control “eyeballs” may view neighboring autonomous systems, as well as remote autonomous systems offering new services, as sources of revenue. This has the potential to create the untenable condition in which a new service might have to pay every “eyeball” autonomous system in the world in order to reach its customers, rather than the current model in which a new service simply pays its transit provider. In the USA, the FCC tried to address these issues with the Open Internet Order.²⁶ However, this has recently been struck down by the Washington, D.C. Circuit Court of Appeals, which contends that the FCC lacked the authority to implement the rules in its Order.²⁷

The inter-domain routing system is not merely technical infrastructure for interconnection, but is also a critical site at which economic disputes play out within the infrastructure of the Internet. While such disputes are to a degree part of normal dynamics in the market for interconnection, these disputes have the potential to degrade, or even partition, the global Internet.

4.2.5 IP Address Space Exhaustion

The gradual exhaustion of IP address space has created a distinct set of concerns for autonomous systems in recent years. Lacking adequately sized IP address prefixes, existing autonomous systems are limited in their ability to expand their infrastructure. These concerns were anticipated as early as 1994, when Network Address Translation (NAT) was developed at the IETF to provide the means for a network to be hidden behind a single IP address (Egevang and Francis 1994). NAT is commonly used in small office and home networks, to allow multiple computers to connect to the Internet through a single IP address provided by an ISP. As the amount of IPv4 address space available has diminished, carrier-grade NAT systems have been developed to allow large ISPs to service multiple customers with a single IP address. However, even these systems are subject to limits, as the number of connections per IP address is limited to 2^{16} , the number of possible TCP ports. Modern web services, such as Google Maps, make tens of requests simultaneously for a single map, effectively limiting the number of possible customers behind a NAT’d IP address to a few hundred [F-SANOG18:82]. In addition, NATs can impose performance penalties on a variety of services through the additional overhead they impose, which is particularly a problem for large scale implementations, such as carrier-grade NATs (Donley et al. 2013).

²⁶See <http://www.fcc.gov/openinternet>, last retrieved Jan 28, 2014.

²⁷See <http://www.freepress.net/press-release/105543/court-strikes-down-fcc-open-internet-order>, last retrieved Jan 28, 2014.

New autonomous systems may find themselves without the IP address prefixes that they need in order to interconnect with other autonomous systems. The current version of IP deployed on the Internet, IPv4 (IP version 4), offers 2^{32} possible addresses; although, as I have already noted, the number of prefixes available is as much a concern as the number of individual addresses available. A new version of IP, IPv6 (IP version 6), has the capability to ameliorate this problem, but offering 2^{128} possible IP addresses.²⁸ Since IPv6 is not backward compatible with IPv4, autonomous systems have to put in place new hardware and software, and develop new skills amongst their network administrators, in order to manage their adoption of IPv6. More importantly, a new IPv6 Internet will only emerge as autonomous systems which adopt IPv6 establish interconnections with other IPv6-capable autonomous systems. Even with the exhaustion of IPv4 address space as a driver for change, the transition to IPv6 has been slow, because of the costs involved in transitioning, and because of the limited benefits, given the relatively small number of IPv6-capable autonomous systems to interconnect with. The transition to an all-IPv6 Internet requires global coordination, similar to that involved in the transition to TCP/IP from NCP described in chapter 3, but of exponentially greater scale and complexity. In the short to medium term, it is likely that autonomous systems will have to offer IPv4 services, potentially “dual-stacked” with IPv6 services.

If IPv4 space available for allocation is limited, how are autonomous systems to manage planning for their infrastructure? The organizations involved in the ICANN regime (which I discuss in the next chapter) have responded by modifying their IPv4 allocation policies to manage scarcity. In addition, a number of intermediary organizations have emerged to facilitate the sale of IPv4 address space in an open market. The emergence of these market-making intermediaries has raised two concerns. First, that the price per IPv4 address might rise to the point where only the largest autonomous systems would be able to afford them. A senior member of the South Asian network operations community summed up this concern up well, commenting at a training session that, “a market for IPv4 addresses creates barriers to entry, and condemns the less affluent to the tyranny of NATs” [F-SANOG18:23]. Second, that IPv4 prefixes would be fragmented to the maximum extent possible in order to facilitate their sale in small lots. This would result in further expansion of the number of prefixes in the default-free routing table, increasing the load on BGP routers in the default-free zone (Edelman 2009). The scarcity of IPv4 address space increases risk and uncertainty for individual autonomous systems. It may also affect all autonomous systems which carry full or near-full routing tables if market solutions to IPv4 scarcity result in fragmentation of available IPv4 prefixes.

²⁸In practice, it would be more accurate to say that IPv6 offers 2^{64} possible IP addresses, since a 64-bit prefix is allocated to each host, allowing multiple virtual network interfaces for each physical network interface.

4.3 A Collective Action Problem

As I have illustrated, a wide variety of significant risks and uncertainties are built into the design and operation of the inter-domain routing system. These risks and uncertainties cannot be addressed only through regulation or a market solution. In fact, as I have shown, a market solution can sometimes be the *cause* of risk and uncertainty, rather than a solution to these concerns. Given the complex global nature of the inter-domain routing system, and the notions of autonomy encoded in BGP, it is near impossible for any regulatory authority to effectively police the behavior of the many thousands of autonomous systems which make up the inter-domain routing system.

Risks and uncertainties in inter-domain routing are *transitive*: inaccurate BGP announcements from individual autonomous systems (such as in the Pakistan/YouTube incident) can spread transitively to remote locations, and potentially to all autonomous systems on the Internet. These risks and uncertainties are also *collective*: fragmentation of IP address prefixes by individual autonomous systems has aggregate effects which affect all autonomous systems carrying a default-free, or almost default-free, routing table. These risks and uncertainties are embedded in Internet protocol standards, which in turn help determine the possible institutional structures for the allocation of resources to operate these protocols, and are enacted within the global topology of interconnections.

Even if the protocol standards were to be modified, these modifications would likely not eliminate these risks and uncertainties, but rather have the effect of shifting the responsibility for these risks and uncertainties. For instance, the IETF’s Secure Inter-Domain Routing Working Group (SIDR WG) has developed new extensions to BGP to secure inter-domain routing with the Resource Public Key Infrastructure (RPKI). These extensions allow autonomous systems to validate the routing and origin claims in BGP announcements they receive, relying on centralized authorities in order to perform this validation. Effectively, risks and uncertainties in BGP announcements are shifted from being a distributed topological concern, to being an institutional concern, reliant on the policies of the centralized authorities enabling validation. These changes also imply a shift in the range of possibilities for governance, as “autonomous” systems will have to give up a degree of autonomy in order to enable policing by centralized authorities. RPKI has been slow to be deployed, because of the additional processing requirements for routers to perform this validation, the need for global coordination to make it effective (as with the IPv4 to IPv6 transition), and worries over loss of autonomy.

There is another perspective through which we could view the risks and uncertainties in inter-domain routing: not as insecurities to be remedied, but rather as necessary costs to gain the benefits of autonomy, and the consequent ease of attachment of new autonomous systems and services to the global Internet. In such a system, failures may be ongoing, but as long as they affect only localized segments of topology, not seen as an immediate cause for concern. While localized failures may occur regularly in this kind of arrangement, it is resilient against global failure. In contrast, a model of centralized authority is susceptible to global failure through political capture, or through failures of centralized technical infras-

structure. As I illustrated earlier, market solutions are also susceptible to global failures as they will likely disproportionately privilege certain actors over others. The socio-technical organization of the inter-domain routing system creates conditions in which *topology* is a significant location of power, supported by specialized centralized institutions, rather than directed by them. It is important to note that my definition of topology does not reduce to market organization. Indeed, one of the important questions I seek to address is how business and technical personnel at autonomous systems reconcile economic logic with the technical logic of interconnection.

Given the seriousness of the risks and uncertainties I have outlined in this chapter, it is a wonder that our everyday experience of the Internet is stable. A critical question for Internet governance, then, is *why* and *how* this stability is currently maintained. I argue that risks and uncertainties in the Internet's inter-domain routing infrastructure should be understood as a collective action problem on a massive scale, concerned with maintaining the integrity of a global commons: the default-free routing table. This collective action problem is addressed through a variety of topological arrangements, and institutional structures for resource allocation and standards-setting activity, which I explore in detail in the following chapters.

Chapter 5

Distributed Governance I: Institutional Anchors

5.1 Encountering the Internet

Over the next three chapters, I explore the internal organization of the Internet, towards developing a perspective on distributed governance. My focus continues to be on the social and technological arrangements which allow the Internet to be an *internetwork*, a network of independently administered networks. In my research, these arrangements initially became visible through publicly available email lists, and standards and policy documents of the Internet's technical communities. The moment at which these arrangements truly became clear, though, was at the first meeting of these communities that I attended, the 74th IETF (Internet Engineering Task Force) meeting in San Francisco in March 2009 (IETF74).¹ The IETF does a great deal of work through online tools such as mail lists and shared document repositories. However, attending the IETF in person brought home to me exactly how important physically co-located social interaction is to the process of producing the virtual spaces of the Internet.

The IETF breaks its activities into a series of broad technical areas, under which a number of working groups are formed, each with a well-defined charter.² I had already developed an interest in inter-domain routing when I decided to attend IETF74, which led me to two working groups in particular: the Inter-Domain Routing Working Group (IDR WG), and the Secure Inter-Domain Routing Working Group (SIDR WG). The IDR WG manages the development of BGP and its associated protocol standards, while the SIDR WG focuses on extensions to add layers of security to BGP.

During the SIDR WG meeting, a furious argument broke out amongst the older members present, in the midst of which one of those involved in the argument turned to the room at

¹See Chapter 3 for a discussion of the origins of the IETF.

²For a complete list of IETF areas and working groups, see <http://datatracker.ietf.org/wg/>, last retrieved Dec 17th 2013.

large, smiled widely, and told us not to read too much into the intensity of the argument, exclaiming, “we’re all old friends here”. This was a moment which stayed with me throughout my research, revealing the social world of the Internet: in the ongoing production of the Internet at the IETF, I found grizzled veterans of technology able to engage in vociferous debate with one another, while maintaining and growing their friendships, and establishing mutual respect. Yet at the same time, they consciously invite others to participate in their discussions, defusing the appearance of tension for those observing their argument from the sidelines. This also made clear to me the importance of in-person meetings; no amount of virtual engagement could possibly capture the passion and physical energy of the interactions at this meeting, whether for those involved in the argument, or for me as a participant observer.

Governance is commonly understood as being a centralized function, relying on hierarchical systems of power to enforce authority and order. Yet as a senior member of the Internet’s technical community said to me rather pointedly, “I object to the term Internet governance. If you use that term, you’re already in the wrong space. The Internet is about coordination!” [F-NANOG56-ARIN30:39]. If governance is to be distributed, rather than centralized, what will serve to hold it together, to construct authority and order? In the case of the Internet, I argue that this “hanging together” of distributed governance is enabled by the essential quality of trust, as evinced by my observations at the SIDR WG: “we’re all old friends here.” Yet trust relationships are in themselves not sufficient to provide a stable system of governance. These relationships are premised upon “institutional anchors” which perform certain functions to enable coordination and collaboration through relationships of interpersonal and generalized trust. These institutional anchors are in turn produced through the trust relationships they help sustain, and also act as meeting points at which the distributed system of trust relations in governance comes into conversation with political and economic interests. As meeting points of varied interests, these institutions are ideal locations from which to understand how the principle of acting “for the good of the Internet”, often espoused by the technical communities I studied, is put into practice.

In this chapter, I examine institutions responsible for standards-setting activity - the IETF - and for the management of critical Internet resources - the Internet Corporation for Assigned Names and Numbers (ICANN), and the American Registry for Internet Numbers (ARIN). The account I present here is through the eyes of the Internet’s technical communities, rather than administrative staff from these institutions, or representatives of purely economic or political interests in these institutions. I do, however, bring in administrative, political and economic voices insofar as they are part of the public records that I examined, and in the cases where my interviewees’ identities integrated these concerns.

I did not plan for any of these institutions to be central to my fieldwork. I originally anticipated that I would have to engage with them to a limited degree, only as a backdrop for the study of the distributed system of trust relationships in Internet governance that I was most interested in. However, as my fieldwork progressed, it became apparent that these institutions were of greater salience to distributed governance than I had originally anticipated. Conversations about these institutions were commonplace in my fieldwork, and

I found that several of my interviewees spanned multiple institutional contexts. For instance, I met one of my first interviewees at the IETF SIDR WG meeting I described here; and then again at ARIN meetings; and at NANOG meetings. As I engaged with these institutions, it became clear that there were significant insights to be gained about the distributed system of governance on the Internet through an examination of the interests and ideals at play within these institutions, as well as the social relations connecting these institutions with one another, and with the Internet's technical communities, represented by professional associations of network administrators such as NANOG.

Rather than think of these centralized governance institutions purely in structural terms, I view them in relational terms. How do each of these institutions function as a bundle of social relations internally? How do individuals find their way into participation in these institutions? What are the external tensions represented by the nexus of social relations visible within an institution? How do these institutions relate to one another? What are the historical processes of change through which these institutions have arrived at their current forms? These are amongst the critical questions posed by a relational perspective which I address here.

I analyze the functioning of these centralized governance institutions, and of their relationship with the Internet's technical communities, using perspectives of *embeddedness*. Through their governance of standards and resources, these institutions play a critical role in the managing the availability of resources, and the openness of standards which are critical to the operation of the Internet. In doing so, they each form a nexus of relations connecting technical, economic and political interests. I use embeddedness in two different ways to understand the role that the Internet's technical communities - representative of distributed trust relations - play in these institutions. First, in terms of "embedded autonomy" (Evans 1995) to understand the embeddedness of technical communities in these governance institutions for the purpose of making their interests felt in policy, while at the same time allowing administrative staff in these governance institutions to function autonomously for the common good. Second, in terms of the embeddedness of markets (Polanyi 2001; Granovetter 1985), to understand how markets for - or market interests in - resources and standards are embedded in the social relations of the Internet's technical communities. Embeddedness is critical to understanding the role of distributed trust relations in constructing these centralized governance institutions, and to understanding the role of these institutions in anchoring distributed trust relations by remaining trustworthy in their operations.

5.2 Developing Standards

The IETF is amongst the most unusual standards-setting organization in the world. It is not formally recognized by any international treaty or industry organization, so it lacks the ability to explicitly force or check compliance to its standards. It has no well-defined notion of membership, whether through nomination or fee payment, as most other standards-setting bodies do. It doesn't even have a formal process for setting standards; rather, standards

are set through “rough consensus and running code”.³ All of these characteristics were understandably features - rather than problems - in the limited research-oriented world of the NSFNET, but it is surprising that these characteristics continue to persist in the global commercial Internet. How, then, does the IETF function, and why is it able to continue to maintain its position as the recognized authority for the setting of Internet protocol standards? How are those involved in the IETF able to balance the political and economic interests of the organizations they represent against developing standards for “the good of the Internet”? It is important to address these questions in order to place the IETF within my broader narrative of distributed governance for the Internet. My aim in this section is not to discuss minutiae of the IETF’s structure and process, but rather to consider the IETF as a site of cultural production of - and contestation over - ideals of openness, transparency and trust.

5.2.1 Participation and Influence

It is remarkably easy to attend an IETF meeting, since there are no pre-requisites for registration; attendees need not represent a particular country or organization. This is an intentional articulation of one of the IETF’s core principles, that IETF participants should leave their affiliations at the door, towards the end of developing standards on their technical merits, rather than to follow the needs of particular interest groups. The IETF’s mission statement (Alvestrand 2004) states this principle clearly:

Individuals who participate in the process are the fundamental unit of the IETF organization and the IETF’s work. The IETF has found that the process works best when focused around people, rather than around organizations, companies, governments or interest groups. That is not to say that these other entities are uninteresting - but they are not what constitutes the IETF.

The IETF has no formal membership. To participate in any IETF activity is to be a “member” of the IETF. As codified in a document describing the organizations of the IETF (Hovey and Bradner 1996):

... membership in the IETF and its Working Groups is defined to be established solely and entirely by individual participation in IETF and Working Group activities.

The “Tao of the IETF” (Hoffman 2012), a document intended to introduce newcomers to the idiosyncracies of the IETF’s organization and process, states that the IETF is not a conference, but rather a meeting of “volunteers”, a “collection of happenings”. All of these documents frame participation in the activities of the IETF as a form of service, developing standards for the common good. This spirit of service is illustrated well in the five “core

³See Chapter 3 for the origins of this phrase.

values” of the IETF elucidated at a plenary of the Internet Engineering Steering Group (IESG)⁴ at the 55th IETF meeting (listed in Davies and Hofmann (2004)):

- “Cares for the Internet”
- “Technically Competent”
- “Open Process”
- “Volunteer Core”
- “Rough Consensus and Running Code”

Although the phrase “good of the Internet” is rarely found in IETF communications, the spirit of this phrase is very much part of the first of these core values. This sentiment is echoed in the opening sentence of the IETF’s mission statement: “The goal of the IETF is to make the Internet work better” (Alvestrand 2004). The second core value establishes a minimum bar for engagement in the IETF, that participants must be competent in the areas of technology for which standards are to be developed. The third core value establishes the importance of openness to the IETF process: all standards, and the standards process itself, are open. This openness involves both the ability to participate in the standards process, as well as the public availability of standards documents and documentation of the standards process itself. The fourth core value reiterates a point I made earlier, that participants should think of themselves as volunteers, acting in service of the Internet. The final core value establishes the process through which standards are developed, using David Clark’s famous saying⁵ to remind IETF participants that the IETF eschews the bureaucracy of traditional top-down standards organizations⁶ in favor of consensus amongst participants, and demonstrable implementations of a standard. The modern IETF in many ways aims to carry on the spirit of openness and collaboration that originated amongst the researchers involved with the NWG, which played a similar role in the development of the ARPANET.

This is not to suggest that commercial and political concerns are absent from the IETF. Participants who attend IETF meetings need to have the resources to pay for conference registration fees,⁷ and hotel and travel costs. These costs can be prohibitive, especially considering that IETF meetings occur three times a year, and are held around the world.⁸

⁴The IESG is responsible for managing the standards process at the IETF. It is made up of Area Directors who are nominated to serve on it. IETF Working Groups are organized into broad subject Areas, each of which is assigned one or more Area Directors. The list of IETF Areas and Working Groups may be found at <http://datatracker.ietf.org/wg/>, last retrieved Feb 7, 2014.

⁵See Chapter 3.

⁶Recall that David Clark’s saying opens: “We reject: kings, presidents and voting.”

⁷Currently USD 650. See <https://www.ietf.org/registration/ietf89/ietfreg.py>, last retrieved Feb 4, 2014.

⁸Although the majority of meetings are still held in North America. For records of past IETF meetings, see <http://www.ietf.org/meeting/past.html>, last retrieved Feb 5, 2014.

Regular IETF attendees require some level of organizational support to sponsor their participation. The degree to which the interests of a sponsoring organization can be reconciled with the IETF's "core values" is a question which IETF participants regularly grapple with.

Many of these sponsoring organizations are vendors of networking equipment - such as Cisco Systems and Juniper Networks - which must conform to IETF standards, or ISPs who deploy this equipment. Governments may also be involved in shaping IETF standards through selective grants. For instance, Kuerbis (2011) traces US government funding to organizations involved in sponsoring individuals involved with the IETF working group for standards to secure DNS. In addition, several non-profit organizations involved in Internet infrastructure development are represented at the IETF. These non-profits typically function through funding from networking equipment vendors, ISPs or governments, but may also draw funds from parallel for-profit arms. For example, the Internet Systems Consortium (ISC) which - amongst other things - develops widely-used open source software for DNS and the Network Time Protocol (NTP)⁹ commonly has employees attend the IETF to participate in the development of these protocols.

The "core values" of the IETF listed earlier were framed at a moment of crisis, when many felt that the IETF had been taken over by the commercial interests of network equipment vendors and large ISPs, and that its bureaucracy had become unwieldy. In a scathing critique of this shift, a highly respected IETF participant went so far as to rebrand the IETF as the "IVTF", the Internet *Vendor* Task Force (Bush 2005). Several of my interviewees indicated that the worst of this over-emphasis on vendor interests has since been ameliorated. This was also a period of declining attendance at the IETF, resulting in financial stress on the operations of the IETF, even with increased attendance fees (IAB Advisory Committee 2004:§2.2.1).¹⁰ However, conversations about the unwieldiness of the IETF bureaucracy remain current. An interviewee who has long been involved with inter-domain routing related her perspective on the changes in the IETF:

... at that time [in the NSFNET period] we would just write the document and start to implement it and deploy it. It's not like today, you have to debate, ... if you have a new protocol, its going to be debated for a long time, because there are a lot of commercial interests, companies interests behind it, all these things, right. [I5:3]

The IETF is also limited by the number of "technically competent" people who are able to commit time and effort to the IETF process. In consequence, some working groups are under-resourced, while others may live beyond their time due to the attachment of participants "tinkering with the edges of the protocol" [F-NANOG57:509].¹¹ A conversation over a lunch

⁹A technology for synchronizing time on computers across the Internet

¹⁰It is worth noting that this moment immediately followed the collapse of the "dot-com bubble" in 2001, which unsurprisingly affected those who built Internet infrastructure, just as much as it affected the Silicon Valley software industry.

¹¹IETF working groups are created to address specific problem areas, and are meant to be dissolved once sufficient work has been conducted to address the problems for which they were created.

I had with two regular IETF participants revolved around the continued relevance of a particular IETF working group [F-NANOG57:509-511]. They were thinking about how best to disband that working group, so that those involved with it could be moved to other areas in need of “people with clue”, especially IPv6.

The notion of “clue” is synonymous with technical competence, as recognized by others within the IETF community. “Clue” is also commonly used within the network operations community at NANOG, with individuals characterized as being “clueful”, or taking a “clue bat” to beat sense into those lacking “clue”. Recognition of expertise is a matter of reputation, which must be earned through repeated and regular interactions within these technical communities.

I ran into these issues early in my research, when I conducted a phone interview with a highly respected IETF participant, who is also a regular attendee at NANOG and other technical community meetings. I had been warned by the person who provided me with an introduction to expect a difficult exchange. The interview turned out to be the shortest that I conducted throughout the course of my research. My interviewee at one stage rather pointedly told me to become better informed on the processing requirements for BGP routers, shortly after which we concluded the interview.¹² When I encountered this particular interviewee in person during my fieldwork, I avoided him quite purposefully at first. When I eventually was able to overcome my unease and approach him - after having attended several NANOG meetings, and gotten to know many in the NANOG community - he evaluated me quite differently, introducing me to a colleague by email as being someone “who seems clueful and is working in the policy space” [EMAIL:10-22-2012]. His change in his attitude towards me may have been in part due to my own increased assurance within these technical communities, but also likely due to him seeing me engaging with others whom he knew and respected.

The IETF is open to participation, but as my account illustrates, such participation requires commitments of time and money, and a shared recognition of expertise. The latter issue is well-recognized at the IETF, with recent conversations on an IETF-wide mail list devoted to mentoring¹³ and promoting diversity in leadership.¹⁴ In both of these conversations, participants comment on the importance of mentoring newcomers, while at the same time noting that it can be difficult for senior IETF participants to find time to mentor. There may also be a structural issue at play here: one participant commented that working group chairs and area directors at the newcomers’ breakfast¹⁵ spent more time talking to one

¹²In contrast, all of my other interviewees were quite happy to discuss technology in various levels of detail.

¹³See <https://www.ietf.org/ibin/c5i?mid=6&rid=48&gid=0&k1=933&k3=12174&tid=1363292321>, last retrieved Feb 10, 2014.

¹⁴See <https://www.ietf.org/ibin/c5i?mid=6&rid=48&gid=0&k1=933&k3=12158&tid=1363291991>, last retrieved Feb 10, 2014.

¹⁵The IETF provides a breakfast for first-time attendees to get to know senior IETF participants.

another rather than in finding newcomers to participate in their working groups.¹⁶ Another participant summed these issues up well:

In some senses the IETF is phenomenally open, since anyone with an e-mail address can sign up for the mailing lists and join the fun, but in other ways it's really difficult to break in, partly because the topics can be complex and subtle, partly because of a (not wholly unreasonable) impatience with people who out of ignorance or otherwise want to reopen ancient arguments, or who imagine that the way to define a standard is to create a kitchen sink of everyone's favorite featurettes.¹⁷

Developing the “cluefulness” required to participate in the IETF is as much as a matter of developing individual technical knowledge, as it is of learning the language and history of particular working groups and areas.

The IETF is also faced with a problem of non-participation from members of the Internet's operational community, exemplified by NANOG. In several conversations with network administrators, interviewees told me how they had great respect for the IETF, but viewed them as not having operational knowledge:

... those guys [vendors represented at the IETF] often don't have much operational experience, they come from the research side, the protocol development side, they're very sharp, they know exactly what they're talking about, but it's not operations. [I1:14-15]

This is a recognized issue at both IETF and NANOG, with efforts to present new RFCs at NANOG meetings to engage the operator community in conversation, while also trying to promote a greater crossover of individuals attending both IETF and NANOG. This is not to say that network operators are *not* represented at the IETF. A network administrator at a Tier 1 ISP who I interviewed at IETF74 told me how organizational affiliation balances against individual reputation in discussions at the IETF:

In a community like this, reputation [is important] ...and not only that, but who you work for, right? Because here, there's a clear distinction between the vendors and the people who use the vendor's equipment. Historically, the IETF has been about the vendors, but service providers bring a lot of weight and heft because we spend tons of money, right? Also, beyond all of that, we run the networks. We want these things. There's that part of it and that's less about reputation than about being the nine hundred pound gorilla in the room.

¹⁶See <https://www.ietf.org/ibin/c5i?mid=6&rid=49&gid=0&k1=933&k2=67852&tid=1391109017>, last retrieved Feb 10, 2014.

¹⁷See <https://www.ietf.org/ibin/c5i?mid=6&rid=49&gid=0&k1=933&k2=67716&tid=1392071096>, last retrieved Feb 10, 2014.

That's more economics. Reputation is more about if I think you've been successful in running a large network. Quite honestly, made a ton of money. If you made a lot of money. That may not be as important to people here who are more almost academia in a certain sense, but it definitely carries weight if you've been able to deploy a very large network and been very successful in it and, quite honestly, been a leader in a certain industry. That helps. In terms of, have you successfully built a network, managed it, architected it, designed and operationalized it, yes, that gets people's attention, especially the vendors . . .

. . . It's about getting the vendors to understand that we want something done in a certain way, and we want you all to do it the same way so when I hook up vendor A's box to vendor B's box to vendor C's box, I'm not running through a nightmare of interoperability. That's a very big part of it. . . You also have to drive consensus, I think, between all the vendors to be successful. [I15:6-7]

Reputation and affiliation are cross-cutting perspectives through which we can understand the IETF. Reputation focuses on the ability of the technically capable individual volunteer to participate in the IETF process. Affiliation focuses on the economic and political relationships between organizations which can shape the interests driving particular standards. It is important to note that the affiliation perspective does not reduce the IETF to being viewed as a tussle amongst operators and vendors. Rather, it illustrates how the IETF can act as a channel for operators to communicate their needs to vendors: as a mechanism through which the risks and uncertainties elaborated in the previous chapter are opened up for discussion amongst a more diverse group. This interview also brought to the fore the diversity of functions within a large ISP organization - my interviewee told me that he does not attend NANOG, but was able to bring concerns from colleagues who do, and who are more involved in inter-domain routing, to the IETF.

5.2.2 Rough Consensus

Of the five core values of the IETF listed earlier, "rough consensus and running code" comes closest to setting an agenda for the IETF standards development process. "Running code" is relatively easy to evaluate from a purely functional perspective: does the code work? Does it provide adequate performance measures? Does it inter-operate with other implementations of the same standard? "Running code" is useful for showing that a proposed standard is workable. However, it can also sometimes be an impediment. A regular IETF participant told me that "the engineers feel like their toys are taken away", given the amount of time that the IETF can take to work through a standard. Another frequent IETF participant involved in the discussion worried that "running code" implemented before a standard takes form can itself cause problems. He gave the example of "a couple of guys going away for a weekend, having an idea, and implementing it", then rolling that idea into a shipping product; it was another eighteen months before the IETF even had a use case for their idea [F-NANOG57:513].

There is always a tension between standards development and software development. The latter seeks to roll out new features quickly, while the former aims towards “rough consensus” for these features. So how is “rough consensus” formed at the IETF? As it turns out, the answer can be quite vary significantly, depending on the approach that individual working groups take towards this problem.

Unlike many other standards bodies, IETF meetings are not moments at which standards are formed, but are rather punctuations in ongoing conversations on working group email lists. As the Tao of the IETF puts it:

One fact that confuses many novices is that the face-to-face WG meetings are much less important in the IETF than they are in most other organizations. Any decision made at a face-to-face meeting must also gain consensus on the WG mailing list. There are numerous examples of important decisions made in WG meetings that are later overturned on the mailing list, often because someone who couldn’t attend the meeting pointed out a serious flaw in the logic used to come to the decision. Finally, WG meetings aren’t “drafting sessions”, as they are in some other standards bodies: in the IETF, drafting is done elsewhere (Hoffman 2012:§4.2).

IETF meetings are not only about working on standards, but also intended to foster a sense of fellowship and enable the cross-pollination of ideas between different working groups and areas:

The meetings, held three times a year, are week-long “gatherings of the tribes” whose primary goal is to reinvigorate the WGs to get their tasks done, and whose secondary goal is to promote a fair amount of mixing between the WGs and the Areas (Hoffman 2012:§3).

Working group meetings are held in rooms organized with rows of chairs facing towards a table at the front of the room, at which the chairs of the working group are seated, sometimes along with other working group officers as well. Those involved with editing a draft standard present and synthesize issues related to that draft, following an agenda set by the working group chairs.¹⁸ Microphones are placed in the aisles for participants to engage in conversation about a draft, or other issues related to the working group. Those speaking at the microphones follow the convention of stating their name before speaking. This is for the benefit of the person taking minutes, and also for the “Jabber scribe”, the person writing details of what is happening in the room into a Jabber chat room.¹⁹ The Jabber

¹⁸All RFCs begin as drafts, and are elevated to RFC status only after achieving “rough consensus” in the working group, approval by the IESG, and then finally being published with a RFC number assigned by the RFC Editor.

¹⁹The IETF uses Jabber to support remote participation. Jabber is an instant messaging technology, standardized as the eXtensible Messaging and Presence Protocol (XMPP) at the IETF. See <http://www.jabber.org/>, last retrieved Feb 11, 2014.

scribe, and others in the room, also help to bring comments from the Jabber chat room into the physical space of the meeting room. Logs of Jabber chat rooms, minutes of meetings, agendas, draft standards in progress and other working group documents are all archived and made available via the IETF website.²⁰

If conversation over a draft does not seem to reach consensus, the working groups chairs may decide that further discussion is needed, and move on to the next agenda item. The furious discussion at the SIDR WG which I opened this chapter with was one such case. On the other hand, if it does seem as though a consensus has been arrived at, the working group chairs will call for a vote to establish the degree of consensus. At the IDR WG, I witnessed this as a show of hands in favor of a draft. There was no counting of the number of people in the room, or the number of hands raised; the working group chair simply glanced around the room to make a decision as to whether or not the show of hands represented “rough consensus”.

In other working groups, humming may be used as a means of indicating consensus. Those in favor of a draft hum first, and those against hum next (Hoffman 2012:§4.2). Again, it is left to the working group chairs to decide whether or not “rough consensus” has been achieved. In a conversation I had with two regular IETF participants, one of them mentioned that voting by humming was the accepted practice in her working group. The other participant responded by asking how consensus could be evaluated with humming; surely a show of hands was better as a means to actually measure the degree of consensus? The first participant responded that humming was less confrontational. By humming, participants are able to indicate their opinion without having to make their stand visible to the rest of the room. In addition, she noted that participants may indicate varying levels of support for an issue by changing the volume at which they hum! The other participant in our conversation remained skeptical, being of the opinion that participants should have to show their opinions to the rest of the room, but understanding why humming might be more acceptable in certain working groups [F-NANOG57:619].

“Rough consensus” is the means through which inclusiveness and openness are brought into participation and discussion at the IETF.²¹ It is a symbol and manifestation of the notion that the IETF is made up of individual volunteers - rather than representatives of organizations - collaborating to produce standards. It is also an important reminder that hierarchies and the influence of powerful interests are anathema to the IETF; recall that the framing of “rough consensus and running code” has its origins in a conflict over the influence of the ISO in the IETF process.²² This is not to say that the IETF process is free of organizational interests (as has been a problem in past years), but rather that “rough consensus” can act as a counter-balance to well-resourced organizational interests. This counter-balance is achieved through “rough consensus” both as a matter of process, and as

²⁰For example, all materials pertaining to the SIDR WG are available at <http://tools.ietf.org/wg/sidr/>, last retrieved Feb 11, 2014.

²¹I do not mean to suggest by this that “rough consensus” is a stand-in for measures of diversity (such as gender, race and geography).

²²See Chapter 3.

a matter of culture: as IETF participants value “rough consensus” collectively, they build it into their own individual perspectives and interactions with other IETF participants.

5.2.3 Organizing Openness

The notion that “openness” needs “organizing” seems at first glance to be a contradiction in terms. If the IETF is a collection of individuals who “care for the Internet”, gathering voluntarily to form standards, why would it need a formal organizational structure? The immediate response is an easy one: organization is needed to manage scale and the diversity of specializations. The IETF is organized by a separation of working groups and areas, with designated bodies for overall management (the IESG) and long-term architectural perspectives and oversight (the IAB). There is a second sense of “organizing” which is also required for the IETF to function. This involves having a paid staff, or contractors, to manage the basic infrastructure of the IETF, including negotiating terms with conference venues, setting up wireless networks and providing audio-visual equipment at meetings, and maintaining hardware and software for the IETF website. The IETF Administrative Director (IAD), supported by the IETF Secretariat, provides these services to the IETF.²³

The first sense of “organizing” calls for greater involvement of volunteers, to take on more responsibilities within the structure of the IETF. The second sense of “organizing” calls for capital needed to allow IETF activities to operate. I have already explored the first sense of “organizing”, and demonstrated the principles through which volunteer work is structured at the IETF. Here, I shall be concerned with the second sense of “organizing”, to explore how it is that the IETF can avoid having its priorities shaped by those providing the capital for its activities.

From 1987 until 1992, the IETF was operated as an activity of the non-profit Corporation for National Research Initiatives (CNRI).²⁴ The Internet Society (ISOC) was formed in 1992, as professional society to engage with the social, political and technical issues arising from the growth and evolution of the worldwide Internet. The ISOC became the organizational home of the IETF, assuming legal responsibility for the IETF’s activities, including providing legal protection to keep the RFC document series freely accessible, legal protection for IETF officers, and legal protection for intellectual property rights disputes (Huizer 1996).²⁵ This move was in anticipation of a reduction or eventual elimination of US government funds to support the IETF.²⁶ CNRI continued to provide operational support for the IETF until 1998, at which time the IETF Secretariat was moved to a for-profit CNRI subsidiary, with the intention of making the IETF Secretariat self-supporting (IAB Advisory Committee 2004).

²³See <https://www.ietf.org/secretariat.html>, last retrieved Feb 12, 2014.

²⁴CNRI was founded by Robert Kahn in 1986, following his work as Director of the ARPA’s IPTO, during which period he was responsible for co-inventing the TCP/IP protocols with Vint Cerf.

²⁵Legal responsibility for IETF intellectual property rights was shifted in 2006 from the ISOC to a new entity, the IETF Trust.

²⁶See <http://www.internetsociety.org/internet/what-internet/history-internet/ietf-and-internet-society>, last retrieved Feb 12, 2014.

The IETF Secretariat function is now contracted to Association Management Solutions. The RFC Editor function is contracted to the Information Sciences Institute at the University of Southern California.²⁷

The IETF currently manages the majority of its budget from meeting attendance fees, with the balance contributed by the ISOC.²⁸ The bulk of ISOC’s revenue comes from the Public Interest Registry (PIR), in which ISOC is the sole incorporator. The PIR manages the “.org” top-level domain registry, meaning that the PIR receives revenue for every registration of a domain name ending in “.org”, providing a stable source of recurring revenue. In effect, the ISOC is not dependent on large donors in order to be able to fulfill its mission, essentially using the Internet (the “.org” registry) to fund the development of the Internet. By association, the IETF is similarly unencumbered.

The IETF has also been careful to form clear policies on intellectual property rights held in contributions to the IETF, in order to ensure that IETF standards remain open. These are articulated via the IETF’s “Note Well” notice which participants are reminded of at every IETF meeting.²⁹

5.2.4 Opening Standards

As the IETF transitioned from the NSFNET to the Internet, it was faced with political and commercial pressures, which did become dominant at one period in its history. However, it has managed to remain a respected and trustworthy organization in which standards for Internet protocols are developed, in no small part due to the embeddedness of the Internet’s technical communities in its operation. Throughout its history, the IETF has maintained an open standards development process, producing standards in service of the Internet as a whole, rather than in service of economically or politically dominant actors. I have argued that the IETF managed to maintain this openness through an open process and organizational form, a culture which rejects “kings, presidents and voting”, instead valuing technical expertise, providing clearly delineated intellectual property rights policies favoring the public interest, and ensuring freedom from reliance on large donors for operating capital. It is only through the combination, and ongoing negotiation, of these factors that the IETF is able to maintain the openness of the Internet standards process.

5.3 Managing Resources

As discussed in earlier chapters, the Internet relies on a range of different kinds of numbers for its operation, including IP addresses, autonomous system numbers, and well-known port

²⁷This was also the organizational home of Jon Postel, who served as the original RFC Editor. See Chapter 3 for more information.

²⁸As of 2013, the IETF brought in USD 3,570,907 in revenue, and cost USD 5,006,435 in expenses. ISOC financial reports are available at <http://www.internetsociety.org/who-we-are/organisational-reports/financial-reports>, last retrieved Feb 12, 2014.

²⁹See <http://www.ietf.org/about/note-well.html>, last retrieved Feb 12, 2014.

numbers. The operation of the Internet requires these numbers to be uniquely allocated out of a fixed range. This range is relatively small for autonomous system numbers³⁰ and well-known port numbers,³¹ implying a need for centralized arrangements for the allocation of these numbers. The requirement for unique allocation of IP address numbers is a technical decision, not an administrative one, intended to maintain the scalability of the inter-domain routing system through hierarchical allocation of IP address space. This was most recently affirmed in RFC 7020 (Housely et al. 2013), but was detailed in RFC 2050 in 1996,³² which called for:

Distribution of globally unique Internet addresses in a hierarchical manner, permitting the routing scalability of the addresses. This scalability is necessary to ensure proper operation of Internet routing ... (Hubbard et al. 1996:2)

Domain names, which map human-readable text to IP addresses, are another resource, required to make the Internet more easily accessible to ordinary users. Like IP addresses, domain names must also be unique, but not for any technical reason. Rather, domain names must be unique to ensure that users' expectations are met, always directing them to the same location each time they visit a domain. Accordingly, domain names also call forth a requirement for centralized arrangements for their management. These centralized management arrangements map directly to a hierarchical technical arrangement for the management of names, the Domain Name System (DNS).

Even though names and numbers are both governed under different branches of the same administrative hierarchy, the regimes under which they are managed are quite distinct. There are at least two reasons for this separation, which are important to consider for the analysis which follows. First, numbers are limited to a fixed range, while names are potentially infinite. Second, numbers have no intrinsic value; they gain value only through use. For instance, an allocated IP address block has no value until it is actually routed within the inter-domain routing system. In contrast, names *do* have intrinsic value, linked to the international legal regime of trademarks, which must be taken into account in the governance of names.

In this section, and the one that follows, I will focus on the centralized institutional and organizational arrangements for the management of domain names, and the allocation of IP address space, which pertains to my principal problem of inter-domain routing.

³⁰There were originally 2^{16} possible autonomous system numbers. As with IPv4 address space discussed in Chapter 4, autonomous system numbers are also beginning to grow scarce in their original 2-byte range of 2^{16} possibilities. A new 4-byte range of 2^{32} possible numbers is now available.

³¹These are used to uniquely identify a protocol, such as the HyperText Transfer Protocol (HTTP) for web servers which is designated to operate on port 80. There are 2^{10} possible well-known port numbers. For current assignments, see <http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml>, last retrieved Feb 14, 2014.

³²A series of earlier RFCs also dealt with the allocation of address space (viz. RFC 1174, RFC 1366 and RFC 1466), but these were written in the pre-CIDR era of classful IP address space allocation (see Chapter 3), in contrast to the RFCs cited here which rely on CIDR for hierarchical allocation of IP address space.

The Internet Corporation for Assigned Names and Numbers (ICANN) is the entity tasked with global governance functions for Internet names and numbers. It operates under a “multi-stakeholder” model of governance, in which the different stakeholders with interests in ICANN functions are formally represented on the ICANN board, and provides formal mechanisms to engage with ICANN’s administrative structures. ICANN was created in 1998, as a California non-profit entity, performing its functions under a contract with the United States’ Department of Commerce. The Internet Assigned Numbers Authority (IANA) function performed by Jon Postel was folded into ICANN at this time.³³ All Internet governance processes - including those at the IETF - are often viewed as instances of multi-stakeholder governance. This refers not to formal administrative structures, but rather to the principle that Internet governance is open to different stakeholders, instead of being limited to, or privileging, particular stakeholders (as, for instance, nation states are privileged at the ITU). As I will show in the following sections, viewing multi-stakeholder governance as a principle in itself can obscure the actual relations of power at play in different multi-stakeholder governance institutions. It is perhaps more important to pay attention to the substance of the administrative processes, the nature of resources to be administered, and the balance of different interests.

5.3.1 Managing Names

ICANN is responsible for managing the root name servers for the DNS. These servers provide the means to look up the top-level domains (TLD) for names. These include generic top-level domains (gTLD) such as “.com”, “.org” and “.net”, and country-code top-level domains (ccTLD) such as “.in” (for India), “.cn” (for China) and “.br” (for Brazil). The IANA maintains the root zone which has details of all TLDs, distributed through the root name servers. There were originally thirteen root name servers, labeled A through M. These are now managed by different entities, each of which typically maintains multiple geographically distributed instances of the root name server that they operate.³⁴ The root name servers point to the next level of name servers which resolve names for a particular TLD. For instance, looking up “berkeley.edu” will result in an initial lookup of a root name server to find the IP address of the name server which can resolve any “.edu” address. A second lookup will then be performed on the “.edu” name server to find the IP address of “berkeley.edu”.

This technical infrastructure maps to an administrative structure separated into domain name registries and registrars. An entity which manages the infrastructure for a particular top-level domain, under contract from ICANN, is called a domain name registry. For instance, recall that the ISOC is funded by the operations of the PIR, which is the registry for the “.org” TLD. An entity which manages registrations for names in registries is called a domain name registrar, and must be accredited to ICANN. When any organization or person wishes to register a domain for themselves, they must do so via a registrar.

³³See Chapter 3 for more on Jon Postel and the IANA function.

³⁴See <http://root-servers.org/> for details, last retrieved Feb 14, 2014.

gTLD registries and registrars, along with other parties with interests in gTLD operations, are represented at ICANN via the Generic Names Supporting Organization (GNSO). ccTLD registries are represented at ICANN via the country code Names Supporting Organization (ccNSO). ccTLD registries are represented separately since it is recognized that governments have a role in managing the ccTLD for their countries.

Disputes over registration of names are governed by ICANN's Uniform Domain-Name Dispute Resolution Policy (UDRP) across all gTLDs. Individual TLDs may provide specific dispute resolution policies of their own.³⁵ The UDRP was instituted in 1999 to cover issues which arise when a domain name registration conflicts with the interests of the holder of the trademark in the name registered.³⁶ When a complaint is brought under the UDRP, it is managed by one of several ICANN-approved third party dispute resolution providers, which include the World Intellectual Property Organization (WIPO).³⁷ As ICANN has rolled out the recent process for approving new gTLDs, it subcontracted the operation of the Trademark Clearinghouse,³⁸ in which trademark holders may register their marks to preemptively protect them across all new gTLDs. These represent the technical and policy mechanisms through which the intrinsic value of names - as represented in trademarks - are protected under the ICANN regime.

5.3.2 Managing Numbers

The IANA delegates number resources to Regional Internet Registries (RIR), which in turn manage number resources for the regions for which they are responsible. This arrangement was originally envisioned well before ICANN was created; in fact even before the Internet officially came into existence in 1995. By 1992, it had become apparent that a centrally located IANA function would not be able to manage resources for the NSFNET and other emerging IP networks, especially as international connectivity was established. These requirements were formalized in RFC 1366 (Gerich 1992), which noted that "...registries which are located in distinct geographic areas may be better able to serve the local community in terms of language and local customs". RFC 1366 goes on to specify several principles for a regional registry, including that it should be recognized by networking authorities within its geographic area, and that it should be committed to following the policies set by the IANA.

These principles continue to hold true today, as the RIRs represent a "bottom-up" policy development process (which I discuss later in this section) driven by their member organizations, which are typically ISPs and large stub networks within their regions. There are currently five RIRs: the Africa Network Information Center (AfriNIC) for the African region, the Asia-Pacific Network Information Center (APNIC) for the Asia-Pacific region, the American Registry for Internet Numbers (ARIN) for the North American region, the

³⁵For details, see <http://www.icann.org/en/help/dndr>, last retrieved Feb 20, 2014.

³⁶See <http://www.icann.org/en/help/dndr/udrp/policy>, last retrieved Feb 20, 2014.

³⁷For a complete list of UDRP dispute resolution providers, see <http://www.icann.org/en/help/dndr/udrp/providers>, last retrieved Feb 20, 2014.

³⁸See <http://trademark-clearinghouse.com/>, last retrieved Feb 20, 2014.

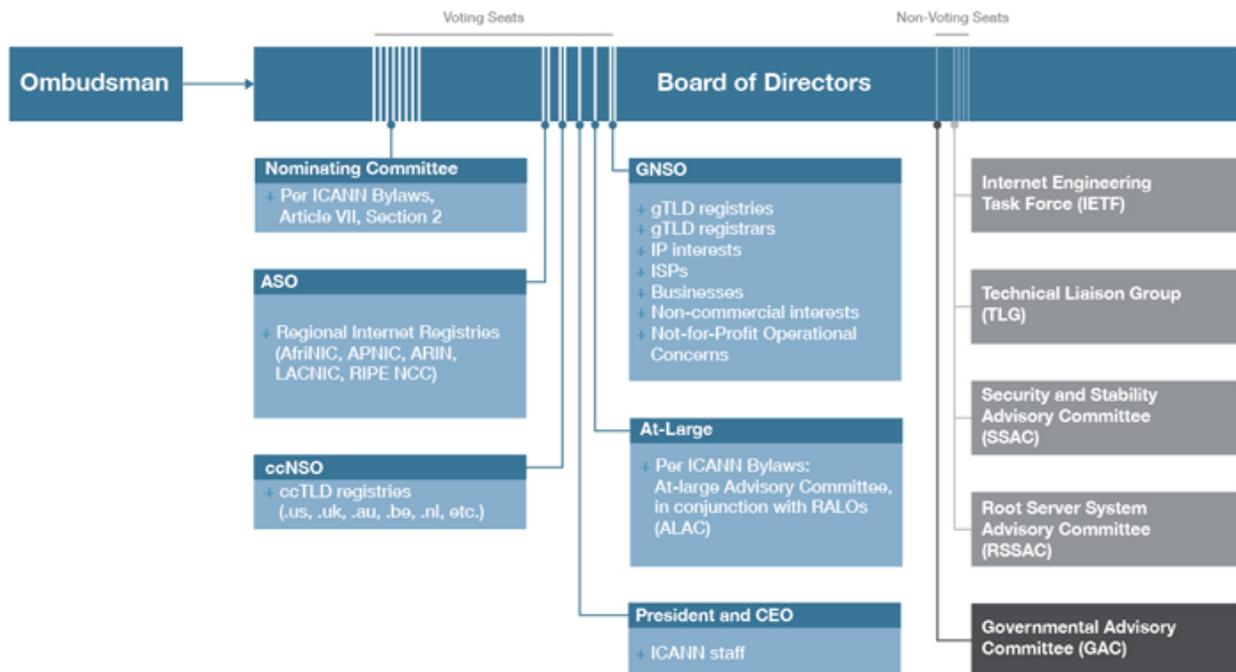


Figure 5.1: ICANN organization chart, available at <http://www.icann.org/en/groups>, last retrieved Feb 14, 2014

Latin America and Caribbean Network Information Center (LACNIC) for the Latin American and Caribbean region, and Réseaux Internet Protocol Européens Network Coordination Center (RIPE NCC) for the European region. Each RIR is free to form policy for number management within its region, which may differ slightly across regions.³⁹

The five RIRs together form the Number Resources Organization (NRO) which advocates on their behalf in a number of policy organizations around the world, including ICANN. The NRO is formally represented in the ICANN administrative structure via the Address Supporting Organization (ASO).

5.3.3 Multi-stakeholder Governance

ICANN provides for public participation in its processes via the At-Large Advisory Committee (ALAC). The ALAC is meant to represent the voice of individual Internet users from around the world at ICANN, and is currently constituted of over 160 organizations.⁴⁰

Each of the groups that I’ve mentioned - the ASO, the ccNSO, the GNSO, and the ALAC - is given voting member positions on the ICANN board, for a total of seven board seats. In

³⁹For a comparison of RIR policies over time, see <http://www.nro.net/policies/rir-comparative-policy-overview>, last retrieved Feb 14, 2014.

⁴⁰See <http://atlarge.icann.org/>, last retrieved Feb 14, 2014.

addition, a Nominating Committee (NomCom), constituted of delegates from these groups and others, is responsible for selecting another eight voting members of the ICANN board.⁴¹ Other groups represented at the NomCom include the Governmental Advisory Committee (GAC) which represents national governments at ICANN, and various specific stakeholder groups (such as the ISP Constituency, and the Non-Commercial Users Constituency). The President of ICANN also serves as a voting board member, making a total of sixteen voting board members.

Additional non-voting board seats are held by board liaisons to the GAC, the Root Server System Advisory Committee (RSSAC),⁴² the Security and Stability Advisory Committee (SSAC),⁴³ the Technical Liaison Group (TLG),⁴⁴ and the IETF.

Voting board members (with the exception of the President of ICANN) are compensated for their time, although they may voluntarily choose to decline this compensation.⁴⁵ Board compensation currently stands at USD 35,000, with an additional USD 5,000 for board committee chairs.⁴⁶ ICANN provides travel support to all board members and liaisons, as well as to selected numbers of participants from various stakeholder groups.⁴⁷

This complicated set of administrative structures, and their relationship to board representation, is illustrated in figure 5.1. In contrast to the “rough consensus” of the IETF, ICANN functions through voting by designated representatives of identified stakeholder groups. The description that I have provided here represents only one level of the complexity of ICANN’s multi-stakeholder governance model, but it is indicative of the range of interests, and the consequent degree of bureaucracy, involved in ICANN’s operations. ICANN operates with close to three hundred paid staff to help manage these interests, and support its mission.⁴⁸

⁴¹See the ICANN By-Laws for details, available at <http://www.icann.org/en/about/governance/bylaws>, last retrieved Feb 14, 2014.

⁴²Which advises ICANN on technical matters related to the root name server system.

⁴³Which advises ICANN on technical matters related to the operation of the Internet’s naming and address allocation systems.

⁴⁴Which advises ICANN on the technical standards which underlie its activities.

⁴⁵As specified in the ICANN bylaws, Article 6 §22, available at <http://www.icann.org/en/about/governance/bylaws>, last retrieved Feb 20, 2014. Compensation for board members was instituted in September 2011; board members served without compensation before this time, and the institution of compensation was a somewhat contentious issue. See http://icannwiki.com/index.php/ICANN_Board#Compensation, last retrieved Feb 20, 2014.

⁴⁶See <http://www.icann.org/en/groups/board/documents/ce/voting-list-09jan14-en.htm>, last retrieved Feb 20, 2014.

⁴⁷See ICANN’s community travel support policy, available at <http://www.icann.org/en/news/in-focus/travel-support>, last retrieved Feb 20, 2014.

⁴⁸See ICANN’s operating plan and budget for 2014, available at <http://www.icann.org/en/about/financials/adopted-opplan-budget-fy14-22aug13-en.pdf>, last retrieved Feb 15, 2014. This lists 306 staff at the end of the year, with 277 staff on average through the year.



Figure 5.2: The view from the audience at an ICANN board meeting, March 18th, 2011

5.3.4 At an ICANN Meeting

Unlike the IETF, registration for ICANN meetings is free. However, ICANN meetings are held three times a year at different locations around the world, requiring significant hotel and airfare costs for regular attendees. In order to engage meaningfully at ICANN, an attendee must align with one or more stakeholder groups. As discussed earlier, many of these groups represent industry bodies, such as registrars and ISPs, or other entities in the Internet governance ecosystem, such as the RIRs or the IETF. The sole exception to this rule is the ALAC, which provides the means for public participation, but is typically composed of representatives of independent organizations. In contrast to the IETF, which in principle rejects organizational affiliation, participation at ICANN requires attendees to represent the specific industry and organizational interests of the stakeholder group with which they are aligned.

ICANN meeting schedules are complex, with multiple stakeholder groups and committees meeting in parallel, and sometimes meeting with one another or the ICANN board for specialized consultations.⁴⁹ I attended the board meeting at the 40th ICANN meeting in San Francisco on March 18th, 2011.

The room for the ICANN board meeting was organized in a manner similar to that for IETF working groups, albeit on a much larger scale. Ranks of chairs faced the front of the room, at which the ICANN board sat on a raised dais, flanked by projector screens (see figure 5.2). Microphones in the aisles provided the means for attendees to engage in conversation with the board. Booths at the back of the room were provided for translators (translating live to Spanish and French) and live transcription, which was projected on the screens to the left of the stage. Documents under discussion were projected on the other screen. Remote participation was enabled via live audio feeds directly from the room, with video feeds and

⁴⁹For example, see the schedule from the 40th ICANN meeting held in San Francisco in March 2011: <http://svsf40.icann.org/node/22237>, last retrieved Feb 15, 2014.

comments from remote participants supported using Adobe Connect.⁵⁰

At the IETF meeting I attended, clothing was relatively informal, with most people wearing t-shirts with jeans or shorts.⁵¹ While some attendees were dressed informally at ICANN, there were also a fair number in suits. This was in line with the overall feel of the meeting, which was much more formal and structured than meetings of any of the other organizations I attended for my fieldwork. I had a strong sense of the meeting being almost scripted, with little space for spontaneous conversation.⁵² This is perhaps unsurprising for an organization at which so many different interests - ranging from corporations to governments - are represented.

The board meeting I attended was significant, since the ICANN board was due to take a vote on the approval of the new “.xxx” sponsored gTLD,⁵³ a contentious issue which had been under discussion for seven years by the time of this board meeting.⁵⁴ The process through which the “.xxx” sponsored gTLD was approved is instructive for the way in which it reveals the balance of interests at play in the management of domain names.

At the time of the meeting I attended, ICANN was still formulating the formal mechanism for the approval of new gTLDs, making this an important issue. In addition, the GAC was opposed to the “.xxx” sponsored gTLD. The GAC board representative stated this position quite forcefully when the GAC met with board prior to the public board meeting:

While there are members which neither endorse nor oppose the introduction of the dot xxx top-level domain, others are emphatically [opposed] from a public policy perspective to the introduction of a dot xxx top-level domain. Furthermore, the GAC would like to inform the ICANN board that an introduction of a dot xxx top-level domain into the root might lead to steps taken by some governments to prohibit access to this TLD. The GAC, therefore, calls the board’s attention to concerns expressed by experts that such steps bear a potential risk or threat to the universal resolvability and stability of the domain name system.⁵⁵

It is important to note that the GAC is only an advisory body at ICANN, with representation on the ICANN board, but no vote. The board is obliged to consult with the GAC, but does not need to follow its advice.

There was additional opposition to the “.xxx” sponsored gTLD from the adult entertainment industry, which feared that it would be ghettoized and more easily censored behind a

⁵⁰A proprietary remote conferencing system provided by Adobe Systems.

⁵¹The Tao of the IETF (Hoffman 2012) notes that “[t]here are those in the IETF who refuse to wear anything other than suits. Fortunately, they are well known (for other reasons) so they are forgiven this particular idiosyncrasy.”

⁵²At the ICANN board consultation with the GAC, statements were spoken of as being “read into the record”.

⁵³A gTLD which has a sponsor representing the interests of a particular community.

⁵⁴For a chronology of these discussions, see <http://www.icann.org/en/news/irp/icm-v-icann/icm-icann-history-21feb10-en.pdf>, last retrieved Feb 15, 2014.

⁵⁵From transcripts of the consultation between the ICANN Board and the GAC, available at <http://svsf40.icann.org/node/22801>, last retrieved Feb 15, 2014.

common gTLD. There were extensive comments against this proposed new gTLD from business people and lawyers involved with the adult entertainment industry during the ICANN Public Forum. They emphatically stated that there was no community support for the “.xxx” gTLD, which is a precondition for a sponsored gTLD; the sponsoring organization must represent the community for that gTLD. One of those speaking proclaimed quite bluntly that “the financial interests of the powerful once again trump the interests of the larger community”.⁵⁶

The ICANN board vote on the “.xxx” sponsored gTLD began with three abstentions due to conflicts of interest. The first was a board member employed with a registrar which might process “.xxx” registrations. The second was from a board member who was an officer at Afilias, which had a minority stake in the ICM Registry, the entity applying to manage the “.xxx” registry; this board member had recused himself from all ICANN discussions pertaining to the “.xxx” sponsored gTLD. The third abstention was from a board member who had been invited to be on the board of i4 - the organization sponsoring the “.xxx” gTLD - and had decided to accept this invitation. A fourth board member - ICANN’s then-CEO and president - abstained as well with no reason stated at the time, indicating that he would later provide written remarks to justify his position. It is a tribute to the integrity of the ICANN board that board members with conflicts of interest would admit them publicly, and refrain from voting. However, it is worth noting that the result was that a full quarter of the board abstained or recused themselves from voting on this highly contentious issue.

The board members who were against the “.xxx” sponsored gTLD were opposed for a number of reasons. One board member felt a need to respect the GAC position, worrying that opposing the GAC might further efforts from governments around the world to move control of ICANN functions to other international treaty bodies. Another board member was of the opinion that the sponsoring organization for “.xxx” was in fact not representative of the adult entertainment community. The issue of free speech was also raised, with the worry that this might be used as an excuse by some governments to increase the censorship of the Internet in their countries, potentially leading towards the fragmentation of the globally unique DNS root managed by ICANN. Those who were for this “.xxx” sponsored gTLD largely spoke of this as a matter of process; that ICANN needed to show the global community that its processes were predictable, and that ICANN would have to face these issues again soon anyway. These were especially important considerations at the time, since ICANN was then working on finalizing the process for proposing new gTLDs. The non-voting board representative from the RSSAC commented that while blocking or filtering of specific TLDs is generally undesirable, there was no technical evidence to suggest that these actions in relation to the “.xxx” sponsored gTLD would have any effects different from blocking or filtering of other TLDs.

Eventually, of the board members who did not abstain, three were opposed, and nine were in favor, carrying the motion in favor of a new “.xxx” sponsored gTLD. The room - which

⁵⁶From the transcript of the ICANN Public Forum, available at <http://svsf40.icann.org/node/22229>, last retrieved Feb 15, 2014.

was crowded for this vote - erupted in applause, and emptied rapidly thereafter, leaving a substantially smaller number of people to witness the board's remaining business.⁵⁷

The case of the “.xxx” sponsored gTLD illustrates the potential fault lines and tensions amongst the interests represented at ICANN in this debate: governments (who do not have a vote), businesses interested in operating new gTLDs, and businesses affected by a specific new gTLD. It is critical to consider the ties of board members to these interests, the flows of capital involved, and the symbolic value of particular gTLDs.

5.3.5 Questioning ICANN's Legitimacy

For all that ICANN is constituted as a “multi-stakeholder” model of governance, economic interests at ICANN can outweigh the interests of the Internet's technical communities and nation states, creating conditions under which the ICANN's legitimacy as the root authority for Internet names and numbers is questioned. The problem, I argue, is one of embeddedness: of markets in social relations, and of administrative structures in the constituencies they serve.

Since the ICANN meeting which I attended in 2011, the process for new gTLDs has been finalized, and the operation of many new gTLDs have already been delegated to registries.⁵⁸ ICANN has gained substantial revenue from the new gTLDs, supplementing its revenue from existing registries and registrars. In 2013, ICANN made over USD 39 million from registries, USD 33 million from registrars and USD 158 million in new gTLD fees. ICANN's total 2013 revenues were over USD 233 million, and total expenses were over USD 150 million.⁵⁹ Registries and registrars paying ICANN's fees are, no doubt, making more money themselves. In short, there is substantial profit to be had in the registration and management of names in the DNS.

As I argued earlier, names have intrinsic value, both in terms of capital and in terms of culture. The “.xxx” sponsored gTLD debate illustrated these issues quite clearly. Alongside profit, there is substantial power associated with the operation of the DNS root. There are two distinct sets of challenges to ICANN's legitimacy to wield this power: one from the Internet's technical community, who believe that ICANN no longer represents their interests, and another from governments who believe they should have a greater say in ICANN's operations.

Several of my interviewees - who are active at the IETF and NANOG - told me that ICANN had been “utterly captured by business interests” [F-NANOG57:215], echoing the voices at the Public Forum during the ICANN meeting I attended. One my interviewees recalled disembarking from a plane behind an ICANN board member, and overhearing someone

⁵⁷Data for this account was drawn from my fieldnotes and transcripts of this ICANN board meeting, available at <http://svsf40.icann.org/node/22597>, last retrieved Feb 15, 2014.

⁵⁸For the current list of new gTLDs, see <http://newgtlds.icann.org/en/program-status/delegated-strings>, last retrieved Feb 15, 2014.

⁵⁹From ICANN's 2013 Annual Report, available at <http://www.icann.org/en/about/annual-report/annual-report-2013-en.pdf>, last retrieved Feb 15, 2014.

ask the board member what it was that ICANN did. The board member responded, “We run the Internet”. The problem, my interviewee continued, was that the board member actually believed that statement, as do many at ICANN [F-NANOG57:219]. In another conversation, an interviewee talked about how ICANN was meant to serve the Internet community, but in reality has come to regard its own survival as paramount, making money off markets for Internet resources, rather than serving the interests of the community [F-SANOG18:74]. Commenting on the “.xxx” sponsored gTLD, an interviewee who works extensively with DNS told me:

It’s become an extortion model saying you don’t wanna be associated as conventofthelittlenuns.xxx? Pay ten bucks. That is an extortion model in my world because there’s a moral association. If it was .co, Columbia or if it was .fr - well heck yeah. I don’t really care if I go on the conventofthelittleangels.fr [sic] but .xxx, somebody hates us, they’re gonna register it. Oh, I’m gonna protect you. That’s like the guys that come in the fedora hat and the pin-stripe suit saying, “Hey, that’s a nice shop there. Something bad’s gonna happen to it if you don’t pay me.” That’s extortion. [I28:14]

Regardless of the objective validity of this statement, it does make clear the dynamics and dilemmas of the business of domain names, and the strength of feeling which these evince amongst the Internet’s technical communities. With power over the DNS root comes the ability to arbitrarily expand markets for names by creating new gTLDs. The intrinsic value of names, and the influence of the domain name industry in ICANN stakeholder groups, are creating conditions in which the interests of capital are being viewed by the Internet’s technical communities as a contradiction against the principles of working “for the good of the Internet”.

In contrast to their dim perspective on ICANN and the domain name industry, my interviewees generally held good opinions of the RIRs. It is worth considering how and why the RIRs - as part of the ICANN regime - are insulated from the kinds of issues that I’ve discussed so far, especially since they rely on ICANN’s IANA function for their number resources. The RIRs operate through a bottom-up policy development process (which I discuss in the following sections) in which all of their members may participate. RIR members are typically ISPs and large stub networks. Their economic interests are in using number resources for interconnection, rather than in number resources themselves. Further, the number resources administered by the RIRs have no intrinsic value, and are limited in quantity, resulting in distinctly different governance dynamics, since there is no profit to be made from numbers. There have been some issues as available IPv4 space has been exhausted, but these have not changed the principles under which the RIRs function, as I will discuss later. In addition, the ASO-NRO arrangement acts as a sort of firewall between the bottom-up policy development process of the RIRs, and ICANN. The RIRs are not simply entities which are directed by ICANN policy, but rather are able to present their perspectives to policy bodies around the world (through the NRO), and negotiate policy with ICANN (through the ASO).

The push for greater government involvement in ICANN's affairs initially manifested in suggestions that ICANN functions should be brought under the International Telecommunications Union (ITU). There have been several attempts to make this happen, beginning with statements from the United Nations' World Summits on the Information Society (WSIS), first held in Geneva in 2003, and then in Tunis in 2005. Of the 122 articles in the WSIS "Tunis Agenda for the Information Society", 54 articles fall under the heading of "Internet Governance". These articles cover areas ranging from the control of Internet resources to the economics of network interconnection arrangements.⁶⁰ The issue of government involvement in ICANN came up again most recently at the World Conference on International Telecommunications (WCIT) in 2012. The WCIT was convened to discuss amendments to the International Telecommunications Regulations (ITRs), which were developed in 1988, prior to the creation of the Internet, and had not be revised since then. The treaty for the ITRs is under the authority of the ITU. The ITRs govern the flow of telecommunications traffic across borders, setting rules for managing traffic flows, billing and settlement, the quality of international services, and other issues. Article 9 of the ITRs provided for "special arrangements" which allow for the exchange of traffic falling outside the provisions of the ITRs. This became the legal mechanism by which IP traffic was exchanged on international routes (Internet Society 2011).

There were numerous efforts at the WCIT to bring the Internet under the purview of the ITRs, covering issues from naming and numbering to the interconnection of networks. All of these discussions happened behind closed doors, as is the norm at the ITU. Information about proposals being floated at the WCIT only became available as leaks.⁶¹ Some of the WCIT proposals understandably alarmed many in the Internet's technical and business communities, as well as the US government, which was faced with the possibility of losing control of the root authority for names and numbers. In addition, different countries found themselves in disagreement over different issues. Eventually, many countries, including the US, refused to sign the final treaty text.

However, this has had the effect of pushing ICANN to more seriously consider how to engage with the governments of the world. This process has been accelerated by Edward Snowden's NSA revelations, with a "Global Multistakeholder Meeting on the Future of Internet Governance" to be held in São Paulo, Brazil, in April 2014, involving representatives from ICANN's stakeholder groups alongside ministerial-level representation from countries and representatives of international organizations.⁶² Most recently, the European Commission has issued a statement offering itself as an "honest broker" in negotiations on the future of Internet governance. The European Commission's statement calls for the "globalisation of ICANN and IANA functions", but at the same time affirms ICANN's multi-stakeholder

⁶⁰For the text of the Tunis Agenda, see <http://www.itu.int/wsis/docs2/tunis/off/6rev1.html>, last retrieved Feb 17, 2014.

⁶¹Which were aggregated at <http://wcitleaks.org/>, last retrieved Feb 15, 2014.

⁶²See <http://www.icann.org/en/news/announcements/announcement-11jan14-en.htm>, last retrieved Feb 15, 2014.

governance model, indicating that “top-down approaches are not the right answer”.⁶³ These moves can be understood in terms of governments seeking greater control in ICANN for public policy purposes; but they must also be understood as a response to the perspective that the ICANN regime is an instrument of US hegemonic power. For example, in conversations with individuals involved in global policy spaces such as the Internet Governance Forum (IGF)⁶⁴ and the WSIS, I was told how the secretary-general of the WSIS spoke of Internet governance as being about ICANN, IANA and the contract between the US government and Verisign for operation of the “.com” registry [F-NANOG56-ARIN30:833].

On the one hand, ICANN is faced with governments who seek representation in, or control over, its multi-stakeholder governance model. Until recent efforts, these moves were rebuffed over fears of “top-down” control by governments or international institutions (such as the ITU). On the other hand, ICANN is faced with a loss of legitimacy in the eyes of the Internet’s technical communities, who view it as being controlled by business interests. Both these challenges may be viewed as problems associated with moves to disembeddedness: of markets for domain names from social relations represented by technical communities and states, with associated restrictions on the participation of technical communities and states in ICANN’s “multi-stakeholder” governance model. ICANN - and the broader system of Internet governance - is at a crucial moment in its history, as it attempts to resolve numerous political and economic tensions. The outcomes are far from certain, but its possible futures present a critical set of questions for the governance of the global Internet, and for the design of multi-stakeholder international institutions.

5.4 Numbering Networks

As I’ve already discussed, a range of different kinds of numbers are required for the operation of the Internet. In this section, I focus on the IP addresses required for the operation of the inter-domain routing system.⁶⁵ I will examine the general principles and issues involved in the governance of IP addresses, and the application of these principles within the context of ARIN, the North American RIR. I will show how the distinctive notions of value attached to numbers, and the principles and structures involved in their management, create a regime that is able to maintain the embeddedness of markets for number resources, and function through embedded autonomy, able to serve the common good while embedded in relations with the Internet’s technical community.

⁶³See http://europa.eu/rapid/press-release_IP-14-142_en.htm, last retrieved Feb 15, 2014.

⁶⁴An annual forum at which Internet governance issues are discussed. Attendees include representatives from governments, Internet governance organizations (such as ICANN and the NRO), and civil society.

⁶⁵Autonomous system numbers (ASNs) are also a critical resource for the inter-domain routing system, but are of less concern than IP addresses. There are still adequate numbers of 2-byte ASNs available, and the transition to 4-byte ASNs has been a relatively painless process. It is also much more difficult to spoof an ASN than it is an IP address block.

I draw extensively from the ARIN Number Resources Policy Manual (NRPM)⁶⁶ and the ARIN Policy Development Process (PDP)⁶⁷ for this section. The NRPM specifies the policies for the allocation and management of ARIN’s number resources. The PDP specifies the process through which the NRPM is amended.

5.4.1 Stewardship

The allocation of IP address space is governed by the principle of *stewardship*. The RIRs assert clearly that IP addresses are not property to be bought and sold, but rather resources held in common which are administered by the RIR system.⁶⁸ This assertion is intended allow the RIRs to steward the allocation of available IP address space, to control prices for address space, and to help preserve the size of the inter-domain routing system’s default-free routing table.⁶⁹

As stewards of IP address space, the RIRs are able - to the extent possible - to provide contiguous blocks of address space to requesting autonomous systems, rather than multiple fragmented blocks. The expectation is that autonomous systems receiving a contiguous block of IP address space will advertise it to the inter-domain routing system as a contiguous block, resulting in a single entry to the default-free routing table. In contrast, fragmented address block assignments will result in one entry in the default-free routing table for each fragment. Recall from Chapter 4 that there are 489,912 prefixes currently visible in the default-free routing table, which, if maximally aggregated, would condense to 275,505 prefixes.⁷⁰ There is nothing preventing autonomous systems from fragmenting their allocations when announcing them via BGP, but aggregation is generally considered to be a good practice.

Address space allocations take two forms: provider-aggregatable (PA) address space, and provider-independent (PI) address space. PA address space is allocated to ISPs who may in turn assign fragments of that address space to their customers, with the expectation that they will announce the aggregated address space to their neighboring autonomous systems. PI address space is allocated to stub autonomous systems which are multi-homed, with multiple network providers for redundancy. In this case, a stub autonomous system should announce its PI space as a single block to all of its network providers. Autonomous systems are generally discouraged from applying for PI address space, and are instead expected to obtain assignments from their ISP, or their ISP’s upstream network provider, in service of the principle of aggregation. There are, of course, exceptions to these rules. For instance, a customer which has received a fragment of PA address space from an ISP may also announce

⁶⁶Available at <https://www.arin.net/policy/nrpm.html>, last retrieved Feb 16, 2014.

⁶⁷Available at <https://www.arin.net/policy/pdp.html>, last retrieved Feb 16, 2014.

⁶⁸I encountered this assertion in numerous presentations and materials, and in my interviews. For an example, see this statement from the NRO: <http://www.nro.net/wp-content/uploads/About-NRO-v2013.pdf>, last retrieved Feb 16, 2014.

⁶⁹For example, see the NRPM §1. Similar principles guide the policies at the other RIRs as well.

⁷⁰These figures were accurate at the time that I wrote Chapter 4, and will have changed since then. For current figures, <http://www.cidr-report.org/as2.0/>, last retrieved Feb 16, 2014.

this address space via another ISP; or an ISP may announce different fragments of address space to different neighboring autonomous systems for the purposes of traffic engineering.

ISPs which receive PA address space allocations are called Local Internet Registries (LIRs), and are expected to register any assignments they make from their allocations with their RIR. National Internet Registries (NIRs) administered by governments, may also be created under contract to the RIR for their region. NIRs receive resources from their RIR for distribution within their country.⁷¹ The RIRs receive resources for distribution in their regions from the IANA. This hierarchy of organizations ensures a hierarchy of allocation, in order to maintain global uniqueness of IP address allocations, with regional control over allocation policy.

Following the principle of stewardship, IP address space is allocated by the RIRs through a needs-based assessment process. In North America, entities requesting IP address space must justify to ARIN the amount of space requested, demonstrating efficient utilization of space that they already use (either allocated from ARIN, or assigned from an upstream ISP), and a need for additional space.⁷² These policies are set by the ARIN community, as I will describe in the sections that follow.

5.4.2 Classifying IP Address Space

There are certain special classifications which govern the use of IP address space. Some of these classifications flow from the IETF, some are specified by the RIRs, and some are defined through use - and abuse - within the inter-domain routing system. Note that while all of my examples relate to IPv4 address space, similar considerations also apply for IPv6 address space.

RFC 1918 (Rekhter et al. 1996), RFC 5735 (Cotton and Vegoda 2010), and RFC 6598 (Weil et al. 2012) specify reserved IP address space. This includes private address space for use within a local area network, address space to be used for documentation, multicast address space, shared address space for use in carrier-grade NAT systems, and several other categories of address space. These reserved address blocks are not meant to be announced into the inter-domain routing system; they are called “martians”, since IP packets in the inter-domain routing system originating from these address blocks could clearly not have originated on Earth.⁷³

Blocks of IP address space which have not been allocated by the IANA to any of the RIRs are known as “bogons”, since any announcements of these address blocks, or traffic claiming to originate from these address blocks, must be bogus. Like martians, bogons must be filtered, since they should not be visible in the inter-domain routing system. Martians

⁷¹I will discuss NIRs in the context of India in the next chapter.

⁷²See NRPM §4 for rules governing the allocation of IPv4 address space.

⁷³See the use of this term in documentation from major networking vendors such as Cisco Systems and Juniper Networks: <http://www.cisco.com/cpress/cc/td/cpress/fund/primer/cb0708.htm> and http://www.juniper.net/techpubs/en_US/junos13.2/topics/concept/martian-addresses-understanding.html, last retrieved Feb 18, 2014.

are considered to be included in the set of bogons. As the amount of available IPv4 space in the IANA free pool has reduced, the set of IPv4 bogons has become equivalent to the set of IPv4 martians. A “fullbogon” indicates IP address space which has been allocated to a RIR by the IANA, but not yet allocated by the RIR to one of its members. Fullbogons include the set of bogons and martians, and should similarly be filtered by autonomous systems. The non-profit wing of Team Cymru, an Internet security research firm, offers a range of free services for the benefit of the Internet’s technical communities, amongst which are feeds of bogon and fullbogon lists, provided through a variety of mechanisms, including BGP, entries in RAdB and plain text. They are careful to note that these entries are updated daily as IP address blocks are allocated for use, which means that autonomous systems using their feeds for filtering must update their filters regularly to avoid blocking legitimate network traffic.⁷⁴

In the last chapter, I described the problem that autonomous systems have in evaluating the veracity of the claims made in BGP announcements. The records of IP address space allocations maintained by the RIRs - often in an Internet Route Registry (IRR)⁷⁵ - provide some relief in this regard, although these still rely on ISPs to maintain the records of the assignments that they make to their customers. In the ARIN region, this data is maintained and made available either through a Referral Whois (RWhois) server operated by an organization which received an address block allocation, or through the Shared Whois Project (SWIP), which is a shared database of IP address assignments.⁷⁶ Both of these services allow anyone to look up the details of the entity to which a particular block of address space has been assigned. Since SWIP is a shared database, the data recorded in SWIP is an attractive target for criminals, who may set up entities which resemble legitimate holders of address space in order to steal the use of their address blocks. A network administrator involved in managing IP address space for a Tier 1 autonomous system related these concerns to me:

... a customer will call up and say, can you please route 1.2.3.0/24 for us, and we’re usually pretty diligent about checking the SWIP record, checking out who that address is SWIP’d to, looking at what the customer name is, which through the contracting goes through a credit check, so I have a fairly good belief that our customer name is actually valid. So if our customer name kind of, sort of, looks like the name in SWIP, or if our customer can provide legal documentation showing that their company is also the company who’s listed in SWIP because of a buy-out, or a merger, then we just add it to the prefix list. It’s not a perfect system, not everybody checks. We have seen some very interesting near misses, like “Banks of America” looks an awful lot like “Bank of America”, you have to react quickly, so there have been a number of cases where people set up companies which have spellings that are very close to companies that are now defunct, to try and steal their address space. [I14:5]

⁷⁴See <http://www.team-cymru.org/Services/Bogons/>, last retrieved Feb 18, 2014.

⁷⁵See chapter 3 for details.

⁷⁶For details, see the ARIN page on reporting reassignment information at <https://www.arin.net/resources/request/reassignments.html>, last retrieved Feb 18, 2014.

The Spamhaus Project is an international non-profit organization which provides services to combat email spam and malware, offering support for the instances in which criminals do manage to bypass the checks and balances offered by SWIP and similar services. Amongst the services provided by the Spamhaus Project are the DROP (Don't Route Or Peer) and EDROP (Extended DROP) lists, which track IP address blocks which are hijacked or leased by professional spam and cyber-crime organizations.⁷⁷ Like the Team Cymru martian, bogon and fullbogon lists, these are intended to be used as filters to help block malicious network traffic, and must be updated regularly to avoid inadvertently filtering legitimate traffic. Even so, creative spammers may find a way around these restrictions: at ARIN30, a presenter talked about how malicious email marketers sometimes rotate through multiple \24 IP address blocks in order to avoid getting tagged as spammers [F-NANOG56-ARIN30:810].

In consequence, when address space is returned to a RIR, it may need to be “delegonized” before it can be reallocated for legitimate use. This is often as simple as holding the address space for a while before it is made available for allocation, providing enough time for the change in status of the address space to be reflected in the Spamhaus lists and other similar lists maintained by private Internet security companies and large autonomous systems. In some cases, this may also require direct reporting and coordination with the entities which maintain these lists in order to ensure that the reallocated address space will be globally routed within the inter-domain routing system.

Complicating the ability of the RIRs to manage address space allocations is the fact that each of the RIRs developed along distinct historical trajectories. As I discussed in Chapter 3, Jon Postel originally played the role of IANA, managing all Internet number allocations, including IP address blocks. This function was subsequently managed through a series of arrangements, leading up to the NSF creating the Internet Network Information Center (InterNIC) in 1993 to manage names, and publish number allocations managed by the IANA. ARIN was constituted on December 22nd, 1997, to manage numbers on behalf of the North American Internet community, rather than under contract to the NSF. The historical changes in authority for numbers continue to have salience for the RIRs today, especially in North America, where the bulk of historical (or “legacy”) IP address allocations were made. All address space allocated after ARIN's inception is governed under ARIN's Registry Services Agreement contract. Address space allocated prior to this date does not fall under ARIN's purview, but may be registered with ARIN under a Legacy Registry Services Agreement contract. However, legacy address space holders are under no obligation to register this space with ARIN. Statements from ARIN officers indicate that as much as 50% of all address space in the North American region is unregistered legacy address space, held by a variety of organizations, including US government agencies [F-NANOG56-ARIN30:470]. Some of this address space has been voluntarily returned to ARIN, such as a \8 address block from Stanford, but much of it has not.

This has become a greater problem in recent years, as the amount of available IPv4 space has grown scarce, resulting in an increase in the perceived economic value of IPv4

⁷⁷See <http://www.spamhaus.org/drop/>, last retrieved Feb 18, 2014.

address space. Legacy address space holders, in essence, are holding a resource for which they have not paid, but has all of a sudden become valuable. In response to IPv4 scarcity, several organizations emerged to create an open market for IP address space. This is in direct contradiction to the RIR principle of stewardship which states that IP addresses are not property to be bought and sold, but rather resources held in common. The tussles between ARIN and IPv4 brokers have proceeded both through debates over the ARIN PDP, and in bankruptcy court, where the legal status of IP address space allocated to bankrupt organizations is decided. At the ARIN meeting I attended, I asked a lawyer involved with ARIN about the IPv4 brokers' business model. His response was angry:

I am sort of the lawyer who's Johnny-on-the-bridge, against the horde of other lawyers who are at this end of the bridge, many of them, who want no restrictions at all. Now, that nihilist position is actually antithetical to the operation [of the Internet] because disaggregation [of address space] would occur. In other words, if you take that to the logical end, the Internet will [dying gasps], right. There are public needs that require that our policies be what they are: conservative, stewardship, preserving the public good. Well, they [IPv4 brokers] don't value the public good. So looking around the room now, you have better dressed people who are [IPv4] brokers and lawyers who are coming here, rather than the propeller-heads who are our usual people. [I39:3.00-4.00]

IPv4 transfer policy was a hot topic during the ARIN30 meeting I attended. Over lunch, ARIN staff had arranged placed signs on tables to be devoted to particular policy issues. The group interested in discussing IPv4 transfer policy was so large that they had to draw multiple tables together. Even so, it was so crowded that I was unable to find a seat to listen to those conversations [F-NANOG56-ARIN30:654].

The courts have ruled in favor of IP address space being subject to ARIN's policies. ARIN policy has been amended to allow the sale of IPv4 address allocations, but only if the buyer can demonstrate need for the address space (in line with ARIN's needs-based assessment policy), and signs a Registry Services Agreement contract with ARIN (Edelman and Ryan 2013). This is viewed as a means to deal with IPv4 address space scarcity, and also to bring legacy address space - when sold - under ARIN's purview.

There is also a special category of address space reserved for critical Internet infrastructure, such as Internet Route Registries (IRRs), DNS root servers and Internet exchange points (IXPs).⁷⁸ All the RIRs maintain policies and reserved blocks of IP address space for critical Internet infrastructure allocations within their regions. Depending on their application, these addresses are sometimes *anycast*, announced via multiple geographically dis-

⁷⁸An IXP is a location at which multiple networks interconnect. Many IXPs are operated as non-profit entities by network operator associations, providing benefits to all participants. I will discuss IXPs further in the following chapters.

tributed autonomous systems around the world to point to replicated copies of the servers operating critical infrastructure.⁷⁹

These classifications have material consequences for businesses in need of address space. During debates over a proposed modification to the NRPM to reserve additional IP address space for critical infrastructure, the question arose of whether DNS operators who expected to receive new gTLDs under ICANN's new gTLD process should be allowed to obtain critical infrastructure space for the DNS servers providing new gTLD services.⁸⁰ This position was strongly criticized during discussion over this policy:

While there are some commercially operated Internet exchange points, those are a tiny minority of the total, and the primary beneficiaries of Internet exchange points are the Internet community, whereas the primary beneficiaries of many of the new gTLDs, particularly the brand TLDs, are specific for-profit corporations that are very small constituencies by comparison.

... we have a number of for-profit entities that are paying very large amounts for TLDs and could acquire the address space needed to support those ... In other words, if you're going to - if you need a 24 to run the new TLD, you can get it on the transfer market.⁸¹ That's not going to be an onerous burden.⁸²

The final version of this policy did not make any critical infrastructure allowances for new gTLD operators. These voices - and those who spoke about the tussle over IPv4 trading - echo the critiques of ICANN which I discussed earlier, with a line drawn between efforts performed in the interests of the "Internet community", and efforts in support of profit which run counter to the interests of the "Internet community".

As this account illustrates, the categories for IP address space management rely on a notion of address space as having value only insofar as it is *routed*: announced into, and accepted within, the inter-domain routing system. The decision to route, or not route, is based in part upon the records of recognized Internet governance organizations such as ARIN, and in part upon services provided by non-profit organizations to the Internet's technical communities, who may individually choose how - or if - to use these services. For reserved address space, value is also conditioned upon the use to which it is put, as in the case of critical Internet infrastructure. Following the guiding principle of stewardship, address space is allocated to be *used* efficiently and appropriately, not to be held or traded. It is, therefore,

⁷⁹For instance, the L root server has 146 geographically distributed instances around the world. For details, see <http://root-servers.org/>, last retrieved Feb 16, 2014.

⁸⁰See https://www.arin.net/policy/proposals/2012_6.html, last retrieved Feb 17, 2014.

⁸¹A \24 IP address prefix, which is the smallest size prefix generally accepted for announcement into the inter-domain routing system. The transfer market is a reference to the market for transfer of IPv4 address space, which I discuss in the following sections.

⁸²From my fieldnotes [F-NANOG56-ARIN30:440-457] and transcripts of the discussion on the policy proposal, available at https://www.arin.net/participate/meetings/reports/ARIN_XXX/ppm1_transcript.html#anchor_6, last retrieved Feb 17, 2014.

important to consider how classifications of IP address space are not neutral, but are rather patterned with political and economic considerations.

5.4.3 Organizing ARIN

ARIN is incorporated as a non-profit entity, headquartered in Herndon, Virginia. It obtains the bulk of its revenue from recurring annual fees for the allocation of IP address space and autonomous system numbers in the North American region. Its costs closely track revenue; in 2012 (the latest year for which audited financial information is available), ARIN had close to 15 million USD in revenue, and slightly over 15.5 million USD in expenses. ARIN made up the balance between these figures through revenue from investments.⁸³ ARIN's 2014 budget estimates slightly under USD 17 million in expenses, and a little over USD 17 million in revenue, with a target of 60 full-time employees. ARIN's 2014 budget expenses include slightly over USD 200,000 due to ICANN, and a similar amount due to the NRO; unlike names, there is little money to be had in numbers.⁸⁴ That said, there is money to be had in the IPv4 transfer market, with which ARIN has to remain constantly engaged. The lawyer I spoke with told me how his firm normally paid an annual retainer of USD 60,000 by ARIN. With the emergence of IPv4 brokers, his firm is representing ARIN in court proceedings across the US almost every day, increasing ARIN's legal fees to close to USD 500,000 [I39:5.10].⁸⁵

The lack of profits from numbers is very much by intention. ARIN - like the other RIRs - follows the principle of stewardship to ensure that number resources are available to all who can demonstrate a need for them. The policies to back up this principle have been modified carefully to manage availability in the face of the exhaustion of IPv4 address space. However, while IPv4 scarcity is an issue for address space allocation requests, the cost of addresses is not, by design. The smallest "XX-Small" allocation from ARIN (smaller than a /22 IPv4 address block) costs just USD 500 per year; the largest "XX-Large" allocation (larger than a /12 IPv4 address block) costs USD 32,000 per year.⁸⁶ As an attendee at an ARIN meeting put it to me, these amounts are "rounding errors" [I39:8.45] on most organizations' balance sheets.

ARIN's fee schedule is set by the ARIN board, with public input offered through the ARIN Consultation and Suggestion Process (ACSP).⁸⁷ The ARIN board has seven seats: six are elected by the ARIN membership, and the seventh is the current ARIN president. The board is responsible for approving all changes to ARIN policy. Changes to policy are

⁸³From ARIN's audited financial information for 2012, available at https://www.arin.net/about_us/corp_docs/annual/2012_audited_financials.pdf, last retrieved Feb 19, 2014.

⁸⁴See https://www.arin.net/about_us/corp_docs/budget.html, last retrieved Feb 19, 2014.

⁸⁵These costs are estimated at USD 420,000 in ARIN's 2014 budget, and were estimated at USD 470,000 in ARIN's 2013 budget, available at https://www.arin.net/about_us/corp_docs/budget2013.html, last retrieved Feb 19, 2014.

⁸⁶From ARIN's fee schedule, available at https://www.arin.net/fees/fee_schedule.html, last retrieved Feb 19, 2014.

⁸⁷See <https://www.arin.net/participate/acsp/acsp.html>, last retrieved Feb 19, 2014.

managed by the ARIN Advisory Council (ARIN AC), which consists of fifteen members, all elected by the ARIN membership. Any organization which has a valid ARIN registration services agreement (regular or legacy) is automatically an ARIN member, and may vote in ARIN elections through their designated member representative. All ARIN board members are also ARIN members.⁸⁸ Individuals elected to the ARIN board or the ARIN AC need not be part of ARIN's membership, although they must be nominated by an ARIN member. The only criterion of importance in elections to the ARIN board and the ARIN AC is that candidates have a demonstrated interest and background in ARIN's mission and functions.

All elected ARIN board and ARIN AC seats are volunteer positions without compensation, a point which ARIN staff are careful to point out during ARIN meetings [F-NANOG56-ARIN30:896-898]. ARIN does provide travel and lodging costs to board and AC members for its public meetings and other necessary meetings, if their employers cannot cover these costs.⁸⁹

ARIN holds two public policy meetings every year at locations across the North American region, so anyone who participates regularly in ARIN will typically need an organizational sponsor. In recent years, ARIN and NANOG have coordinated schedules to ensure that at least one of these meetings is held back-to-back with a NANOG meeting. This move, along with a new ARIN public policy consultation session at all NANOG meetings, is intended to increase participation from the NANOG community in ARIN's policy processes. ARIN used to have a registration fee for non-members to attend, but these were eliminated in the interests of encouraging wider participation [F-NANOG56-ARIN30:891-892].

Policy proposals may be submitted by anyone who is not an ARIN board member, or on the ARIN staff; which is to say that *anyone*, not necessarily a representative of an ARIN member organization, can suggest modifications to ARIN policy. New policy proposals are posted to the ARIN website, and assigned a member of the ARIN AC as a shepherd, who works with the policy originator to ensure that the text of the proposal is as unambiguous as possible. The ARIN AC then evaluates the policy to ensure that it is within the scope of the PDP, applies within the ARIN region, and contains clear text proposing modifications to the NRPM. The ARIN AC evaluation, along with the text of the policy proposal, is then posted to the ARIN Public Policy Mailing List (PPML). If the evaluation is negative, this is a public notice of rejection. If the evaluation is positive, the policy proposal is moved to the status of being a draft policy, for further discussion by the ARIN community, on the PPML and at ARIN meetings. Draft policies are also submitted for a review by ARIN staff and legal personnel, who may raise any process or liability issues related to a policy.

The ARIN AC continues to develop the draft policy, taking into account staff and community comments, until such time as they feel that the draft policy can be recommended for adoption. Support for recommended draft policies is polled amongst the ARIN community at ARIN public policy consultations, which may be part of an ARIN public policy meeting,

⁸⁸See https://www.arin.net/about_us/membership/overview.html, last retrieved Feb 19, 2014.

⁸⁹See ARIN board and ARIN AC travel policies, available at https://www.arin.net/about_us/bot_travel.html and https://www.arin.net/about_us/ac_travel.html, last retrieved Feb 19, 2014.

or held as a special track at a NANOG meeting. Remote participants may participate in ARIN public policy consultations via a Jabber chat room. Unlike the IETF, ARIN meetings are not merely punctuations in a ongoing discussion, but rather moments at which community support for a policy proposal is formally evaluated. Before an ARIN public policy consultation, ARIN staff prepare a variety of meeting materials, including a discussion guide with text of all policies to be discussed at the meeting. Transcripts and recordings of ARIN meetings, along with presentation and notes, are made publicly available after the meeting has concluded.⁹⁰

Recommended draft policies with sufficient community support are eventually advanced to last call status, after which they are sent to the ARIN board for review. The board may yet reject a policy proposal, or send it back to the community for further review. Policies accepted by the board are formally made part of ARIN policy by integrating their text into the NRPM.

ARIN meetings are not just about the formulation of policy. They also feature reports on the overall status of ARIN activities, reports from the IANA, the NRO and other RIRs, and sometimes also reports from the IETF or other organizations on activities relevant to the ARIN community. For instance, the ARIN30 meeting I attended featured a report from the IETF on the transition to IPv6, and reports from other RIRs talking about IPv4 utilization rates in their regions, membership growth, training, and other issues of interest.

5.4.4 Negotiating Policy

ARIN attendees must reconcile a variety of interests and principles: the interests of their employers, the interests of the ARIN membership at large, and the principles of stewardship originally conceived at the IETF, shared by the IANA and other RIRs (Hubbard et al. 1996; Housely et al. 2013). All of these are typically wrapped up and balanced in each attendee's individual perspective, and the debates that emerge around policy are often the result of differing perspectives on how to manage this balance of interests and principles. These are shaped by the needs of particular Internet industries, an aim towards common generally applicable policy, and an idealized sense of working towards the common good of the Internet in the ARIN region and beyond.

ARIN attendees are very aware of these internal conflicts, and will almost always specify the affiliation under which they are stating their views, sometimes speaking as an individual, sometimes as an employee of a company, or sometimes as an officer of the ARIN board or ARIN AC. The statement of affiliation is part of the ARIN process (as it is at other Internet governance institutions as well), which attendees are reminded of at the beginning of sessions, and sometimes through notices on microphones (see figure 5.3). It is generally acknowledged that different priorities follow different stated affiliations.

⁹⁰The meeting materials and records of the ARIN30 meeting are available at https://www.arin.net/participate/meetings/reports/ARIN_XXX, last retrieved Feb 21, 2014.



Figure 5.3: A microphone in the aisles at ARIN30, with a notice reminding attendees to state their name and affiliation before speaking.

In this section, I discuss how these interests and principles are balanced in the practice of policy formation at ARIN, using the case of ARIN proposal number 180 (hereafter Prop-180), “ISP Private Reassignment”,⁹¹ presented during the 30th ARIN meeting (ARIN30). For this account, I draw from the emails detailing this policy proposal (linked below), my fieldnotes [F-NANOG56-ARIN30:700-723], transcripts of the discussion around this proposal,⁹² and an email thread discussing this proposal after the ARIN30 meeting concluded.⁹³

As I noted earlier, ISPs are expected to register assignments from their IP address allocations that they make to their customers. These registrations of assignments are made available publicly via SWIP or RWhois. The purpose of Prop-180 was to allow ISPs to make assignments of IP address space to customers private, rather than publicly visible. Anyone looking up the registration information for an IP address block would see the ISP as the registrant, rather than the customer to which it had been assigned. ARIN was to continue

⁹¹All current, approved and abandoned policy proposals are available at <https://www.arin.net/policy/proposals/>, last retrieved Feb 21, 2014.

⁹²Available at https://www.arin.net/participate/meetings/reports/ARIN_XXX/ppm2_transcript.html#anchor_13, last retrieved Feb 22, 2014.

⁹³See the thread titled “POC privacy” at <http://lists.arin.net/pipermail/arin-ppml/2012-October/thread.html>, last retrieved Feb 22, 2014.

to have access to private records of assignments, for the purpose of evaluating an ISP's utilization of allocated address space. The stated intention of this proposal was in the interests of customer privacy, to protect customers from network attacks that might take advantage of their registered point of contact (POC) information, and also for the cases where customers might outsource the management of their network to their ISP, making the ISP the logical point of contact for technical issues.

Prop-180 was first posted as an informational note to the ARIN PPML on August 9th, 2012. The note provided the text of the proposal, and the number it had been assigned for tracking purposes.⁹⁴ After evaluation by the ARIN AC, this proposal was modified and republished to the ARIN PPML on August 16th, 2012,⁹⁵ before being brought up for discussion at the ARIN30 public policy meeting I attended, from October 24th to 26th, 2012.

ARIN meetings have a single track of sessions, unlike ICANN or the IETF, ensuring that all attendees may attend all sessions. The meeting rooms are laid out similar to those at ICANN and the IETF, with rows of chairs and tables facing the front of the room, where the ARIN AC or ARIN board are seated, along with ARIN staff and others, depending on the session. Microphones are placed in the aisles for attendees to engage in conversation with those at the front of the room, and with one another. The dress code is relatively informal, although as the lawyer I spoke to noted, there were some people in suits, many representing IPv4 brokers. Interactions during sessions at the meeting are not as free-wheeling as those I witnessed at the IETF; nor are they as formal and structured as ICANN. Many of the attendees participating in policy discussions clearly know one another, making reference to their knowledge of one another's established positions, or details of past discussions. This makes for a degree of informality. However, the subject matter under discussion - amendments to policy - requires a precision of language which does lend a certain structure to the discussions.

The text of the Prop-180 seemed uncontroversial to me at first glance, and the justifications offered seemed reasonable. The opening statements about the proposal even noted that there is a similar policy in effect at APNIC. The ARIN AC member shepherding the proposal commented that this was one of the few times that they had brought a proposal to the ARIN community before having it advance to draft proposal status. In this instance, the ARIN AC was sufficiently uncertain about the proposal that they felt it necessary to consult the wider ARIN community before proceeding further. The comments during the discussion made it clear why they felt so, quickly making me aware of my own ignorance.

Although attendees normally queue up to speak at microphones of their own accord, an ARIN AC member requested comments from a particular attendee as soon as the proposal had been presented. The attendee who responded was amongst the few people in suits in the room, and turned out to be an Federal Bureau of Investigation (FBI) agent involved in international law enforcement, attending ARIN30 with one of his colleagues. This was

⁹⁴See <http://lists.arin.net/pipermail/arin-ppml/2012-August/025720.html>, last retrieved Feb 21, 2014.

⁹⁵See <http://lists.arin.net/pipermail/arin-ppml/2012-August/025840.html>, last retrieved Feb 22, 2014.

the only policy proposal discussion in which he participated, suggesting that he and his colleague had come to ARIN expressly for the purpose of offering their opinion on this particular proposal. He related how the FBI uses public registration records as a tool in their investigations, as do many other law enforcements in the US and internationally. Adoption of this policy, he suggested, would result in investigations needing to get additional subpoenas, burdening investigators, as well as ARIN and concerned ISPs.

Before proceeding to further discussion, the President of ARIN offered his opinion that the current system scales well because law enforcement doesn't need to bring their requests to ARIN, since this data is available publicly. He also noted that a subpoena may not always be required; ARIN does comply with requests from law enforcement when there are "life-critical" situations, although those are rare.

The first attendee at the microphone offered comments speaking as himself, opening: "I'm strongly in favor of wrapping this up into a wad and throwing it in the trash and not spending any resources at all of ARIN's time." He suggested that if there did turn out to be community support for this proposal, then the ARIN AC should be careful to run it by personnel at ISPs involved in handling network abuse, to get their opinions before proceeding. Others echoed this concern, arguing that malicious actors could take advantage of private registrations to launch attacks which would be harder to trace. As one attendee noted, "bad guys don't like sunshine and this kind of a thing would insert an extra layer of shade."

Another attendee commented that some organizations contract with multiple upstream ISPs, asking each ISP for an IP address block assignment. Without public records, it would be hard to track this kind of behavior, where organizations leverage private registration to accumulate IP address space from multiple providers. This is a concern in general, but especially problematic in the face of IPv4 address space scarcity. He also invoked the sunshine metaphor: "I think that sunshine is the best disinfectant, and I think preserving transparency is more important than any concerns for privacy for businesses."

A few attendees were in favor of the proposal, saying that "customer data is private and proprietary", and arguing that law enforcement should have to go through proper channels in order to get this data. This was framed both in terms of customer data being private to an ISP, and in terms of an ISP functioning as a shield for their customers' privacy. Another attendee responded that similar proposals had been discussed in the past, with exactly the same justifications, since all ISPs would like to make their customer data private. These earlier proposals had been rejected for reasons already raised, since there are "many classes of customers that would like to make their information private because they're trying to avoid blacklisting while doing bad things." He went on to note that there are tools available which harvest public registration data in order to target a competitor's customers. This attendee, and a few others, stated that they would be willing to consider the proposal, if it were scoped more carefully and narrowly targeted at particular kinds of customers and uses.

A representative of the Japanese National Internet Registry noted that they already support this kind of policy in their Whois database, allowing details like the postal address and name to be hidden, although they require administrative and technical contacts to be

public. She indicated that there was no particular burden from law enforcement, since there was still publicly available contact information, and they did disclose private data to law enforcement when requested through proper processes.

Fears of how this kind of privacy-related policy change might affect ARIN's overall authority were also commented upon, as one attendee invoked the specter of the WCIT:

But my most important point to all of this is we are a self-regulating industry. And my biggest concern is if we use a big hammer and swing privacy one way on this, we can find ourselves going from self-regulated to government-regulated WCIT or whatever the case may be, so that really scares me. So I'd rather be fair than regulated.

Given the number of attendees queuing up to comment, the microphones were closed after a point in the interests of time. Once all comments were concluded, a vote was conducted to establish the community consensus on this proposal. There were 135 people present, both in the room and remote. Of those who voted, 19 were in favor, and 33 against. I noticed that the FBI agents numbered amongst those voting against the proposal, unsurprisingly. It is noteworthy how anyone present can participate in the ARIN policy process, including the FBI agents, who were able to engage in discussion and vote, even though they were not part of ARIN's membership. This openness to participation also allowed me to attend and observe this process easily.

Similar issues to those discussed at ARIN30 were raised on the ARIN PPML in the days following the meeting. Several commenters were emphatic that the tech and abuse POCs needed to be public, with one going so far as to suggest that IP address blocks with unresponsive POCs should be unilaterally revoked by ARIN.⁹⁶ One of these commenters was indignant that the proposal had even been considered:

These ideas of hiding POCs are ridiculous! What is the purpose of a "point of CONTACT" if you cannot use it to CONTACT someone?!?!

I constantly use POCs to try to notify resource owners that their resources (usually a server on their network) have been compromised and are behaving badly (i.e. hosting phishing sites or viruses/trojans). I don't get paid to do it - I do it because it needs to be done. If more obstacles are put in my way (i.e. requiring me to use various web interfaces and log in to get the details I need), I will have less and less time to help out the community.⁹⁷

⁹⁶See <http://lists.arin.net/pipermail/arin-ppml/2012-October/026185.html>, last retrieved Feb 22, 2014.

⁹⁷From <http://lists.arin.net/pipermail/arin-ppml/2012-October/026183.html>, last retrieved Feb 22, 2014.

In this vein, the problem of BGP prefix hijacking⁹⁸ was raised, with a commenter pointing out that he would want to ensure that he is available for anyone to contact, in case they notice prefixes from his address blocks being hijacked.⁹⁹

Emphasizing the shared nature of resources used to construct the Internet, a commenter expressed the opinion that:

Businesses using a *shared* resource such as the *shared* public internet not only have no reason to hide themselves, they have an obligation to be reachable by the members of the community they are sharing the resources with. If they don't want to participate in the shared public resource then they don't need shared public resources in the first place. The simple act of reserving a globally unique address consumes a shared resource and obligates the consumer to be reachable, even if the globally unique address is never routed globally.¹⁰⁰

Some commenters argued that there was a measure of merit to the argument behind the proposal, since POCs are used by ARIN to determine authority to modify registration information. They suggested that the proposal could be modified to make such authorized POCs private, while abuse and network operations center (NOC) POCs remained public.¹⁰¹

Raising the question of regulation again, another commenter pointed out that:

The only alternative to public disclosure and the sort of cooperative policing of the internet as it exists today is to form, and fund, an official "internet police" force that has the legal means, and authority, to investigate all the internet players. This is normally how societies deal with this issue in the physical world. For the current experiment to work, privacy does take a back seat. It remains to be seen how long the current situation will last.¹⁰²

This proposal was eventually abandoned by the ARIN AC, as noted in a posting to the ARIN PPML on October 31st 2012.¹⁰³

The debate over this policy is notable for the range of concerns that it raised, and the strength of the opinions offered, especially from those opposed to it. Those in support of the policy raised concerns of individual and customer privacy, framing privacy in terms of freedom. Supporters also raised commercial concerns, arguing that customer data is proprietary information, not for public disclosure.

⁹⁸See Chapter 4.

⁹⁹See <http://lists.arin.net/pipermail/arin-ppml/2012-October/026202.html>, last retrieved Feb 22, 2014.

¹⁰⁰From <http://lists.arin.net/pipermail/arin-ppml/2012-October/026188.html>, last retrieved Feb 22, 2014.

¹⁰¹See <http://lists.arin.net/pipermail/arin-ppml/2012-October/026184.html> and <http://lists.arin.net/pipermail/arin-ppml/2012-October/026190.html>, last retrieved Feb 22, 2014.

¹⁰²From <http://lists.arin.net/pipermail/arin-ppml/2012-October/026199.html>, last retrieved Feb 22, 2014.

¹⁰³See <http://lists.arin.net/pipermail/arin-ppml/2012-October/026221.html>, last retrieved Feb 22, 2014.

Those against the policy argued for maximal openness. Law enforcement opposed the policy on the grounds that it would make investigations more difficult. Those involved with network operations wanted to be able to notify others of - or be themselves notified of - abuse or hijack of IP address blocks, talking about this in terms of being able to offer aid to their technical community. The problem of resource conservation was noted, with arguments presented that this policy would allow organizations to consume IP address space without any way to track their consumption. This was also framed in terms of the public good, with commenters arguing that the Internet is a shared public resource, which calls for open public accounting of resource use. Finally, some commenters raised the issue of the alternative which might come into being if policies such as this one were allowed: a regulated Internet, controlled by a central institution.

This discussion, and the widespread opposition to this policy, illustrates the centrality of concerns of openness and transparency to the Internet's technical communities. In addition, it reveals the processes through which the Internet's technical communities maintain their embeddedness in the ARIN policy process, and maintain the embeddedness of markets for Internet number resources. These processes are at once enacted in policy development which is open to participation, and in operational processes which make data public for easier coordination and collaboration across organizational boundaries, "for the good of the Internet".

5.5 Opening Resource Management

The processes for participation in the management of critical Internet resources are quite open, both to attendance and observation. However, meaningful participation in these processes is another matter altogether, with distinctly different requirements to participate in the governance of names, and in the governance of numbers. ICANN's multi-stakeholder model requires participants to engage with and act through specific stakeholder groups. In contrast, ARIN's model is closer to that of the IETF, allowing anyone with an interest in the management of numbers to participate in ARIN's process. In both cases - as with the IETF - participants require a certain level of knowledge of the subject matter, although the ability to act on this knowledge is conditioned in very different ways. At ICANN, this depends on position within the ICANN structure. At ARIN, this does depend on relationships and reputation within the ARIN community, but most of all on the salience of a policy proposal to ARIN's mission.

The most significant difference between ICANN and ARIN is in their approach to the value of the resources which they manage. Names are intrinsically valuable, and potentially infinite, while numbers - IP address blocks - lack intrinsic value, and are limited in quantity. IP address blocks gain value only insofar as they are routed on the Internet, and even so, the size of the default-free routing table containing all routed IP address blocks needs to be kept under control.

Accordingly, ICANN functions through market principles, creating policies to expand

the market for names, while at the same time aligning with trademark regimes to preserve the value of names. ARIN, in contrast, functions to protect numbers from the market, creating policies based upon principles of stewardship. The result is a system in which the Internet’s technical communities form very different perspectives on these two anchor institutions: regarding ICANN as having been captured by capital, while viewing ARIN (and other RIRs) as functioning in service to the needs of the Internet’s technical communities.

5.6 Open Institutions, Open Governance

Much of the literature and public commentary on Internet governance frames it in terms of multi-stakeholder governance, whether it be at ICANN, ARIN or the IETF. These accounts are almost entirely structural, focusing almost entirely on the stakeholders involved, and the relations of power between stakeholders (DeNardis and Raymond 2013). Yet as the account I have provided in this chapter indicates, the models of governance in each of the institutions I have described vary along multiple dimensions, so much so that they cannot, I suggest, be viewed as varying instances of the same category of multi-stakeholder governance. The governance models in the institutions I studied are premised upon radically different notions of value, different involvements of capital, different models of participation, different participant identities and interests, and different cultures of interaction.

At the IETF, affiliation - a key marker of membership in a stakeholder group - is overtly denied as a factor in participation. My research shows how reputation and affiliation are both key components of a participant’s ability to engage with the IETF standards process. At ICANN, affiliation to a stakeholder group defines an individual’s ability to participate, and in the process represent the interests of that stakeholder group. At ARIN, anyone may participate in the policy development process, with the ARIN AC acting as a mediating body to shepherd and filter policy proposals. Similarly, the culture of free-wheeling discussion at an IETF meeting - emblemized in “rough consensus and running code” - would never fit in the formal culture of an ICANN meeting. The semi-formal culture of ARIN lies somewhere in between these two extremes.

In spite of the many variations between these different institutions, there are some individuals who participate across these spaces, providing linkages between institutions. These linkages are at once implicit - carrying knowledge and interests between institutions - and explicit - representing one institution in another through formal representation, such as the representation of ARIN via the NRO in ICANN’s ASO stakeholder group. At the same time, many participants are able to disaggregate their identities within any one of these institutional contexts, presenting themselves publicly as themselves, as representatives of corporations, or as citizens of countries.

We must also consider the explicit links between these institutions required to manage their everyday activity. The IETF operates under the umbrella of the ISOC, which gets funding from the PIR’s operation of the “.org” registry, under contract from ICANN. ARIN receives number resources for the North American region from the IANA function at ICANN,

and participates in ICANN as part of the ASO. The IETF is formally represented at ICANN, and brings relevant components of the Internet standards development process to presentations at ARIN. These institutions are all tied together in a complex web of dependencies.

Idealized models of Internet governance institutions - such as multi-stakeholder governance - are essential for us to understand and learn from their operation. However, as I have shown, a purely structural approach can mask the actual social processes involved in the production and reproduction of these institutions.

This research also demonstrates the significant amount of work that goes into producing values of openness and transparency, both as principles of operation for these institutions, and as principles of operation for the larger systems which rely on these institutions to anchor them. Openness and transparency are not a natural state of being, but require organization and resources to put them into effect.

Examining openness and transparency as being part of the construction of “the good of the Internet” illustrates how these values play out very differently across different institutions, and the tensions that arise in protection of “the good of the Internet” from the perspective of the Internet’s technical communities. From this perspective, purely structural models of multi-stakeholder governance are a problem, rather than a solution, as they privilege the interests of stakeholder groups over collectively accepted conceptions of the public good. The dilemma this presents is in the question of how to arrive at a conception of the public good; it may well be that there are conditions under which structural multi-stakeholder models can help arrive at such conceptions.

In this sense, it may be productive to examine these institutions from a perspective of embeddedness. Polanyi (2001) warned against disembedding markets from society. He argued that attempts to commodify land, labor and money - which he termed “fictitious commodities” - would result in social upheaval as the logic of the market is imposed on elements which are constitutive of society, not produced for sale as commodities. Similarly, Granovetter (1985) argues that markets are always embedded in social relations; he advises researchers to analyze markets in terms of the social relations and processes involved in specific markets, rather than assuming a common logic across markets. Finally, I draw from Evans (1995) to analyze these institutions in terms of “embedded autonomy”, to understand the degree to which - and processes through which - the Internet’s technical communities are embedded in the centralized governance institutions of the Internet.

These perspectives offer an interesting way in which to look at Internet governance institutions: not simply as institutional arrangements performing certain specified functions, but as the means through which the Internet’s technical communities maintain, and struggle over, the embeddedness of market relations in their social relations, by maintaining their own embeddedness in these governance institutions. My account, then, is not one of alternatives to markets, but rather of the manner of combination of different logics (economic, social, political, technical) to produce sustainable outcomes in the public interest: a stable global Internet.

The challenges for a system of governance aiming to emulate that of the Internet are twofold. First, a culture which enables openness, coordination and collaboration within

institutional contexts is required, in contrast to a culture which emphasizes secrecy and control. As I've illustrated here, and in earlier chapters, this is no small feat: the culture of the Internet's technical communities emerged through historically contingent processes towards the form it has today. Second, governance institutions need to be willing to limit their activity, leaving a not-insignificant degree of risk and uncertainty to be resolved through coordination and collaboration outside of institutional contexts. This is not simply a matter of letting the market solve these issues, but rather of supporting a cultural context which values cross-organizational coordination and collaboration, embedding market relations in social relations. In the next chapter, I will expand on these issues, examining the ways in which network administrators coordinate and collaborate to order the Internet's inter-domain routing system.

Chapter 6

Distributed Governance II: Ordering Topology

6.1 Seeing Dots

At NANOG56, I made the acquaintance of a young network administrator who was attending his first NANOG meeting. Over lunch on the second day of the meeting, he expressed his frustration at the number of people who had been asking him for information about the company he worked for. He had been asked about how many customers they had, their history of corporate acquisitions, and various other details. “How am I supposed to know all this?”, he asked me [F-NANOG56-ARIN30:230]. Looking at his conference attendee badge - printed with his name, his corporate affiliation and autonomous system number - I noticed that he had affixed a green dot to it. At NANOG meetings, colored dots are used to indicate certain roles and functions: black for people working in network security, red for those willing to sign PGP keys,¹ and green for peering coordinators.² Stickers with different colored dots are freely available at the registration desk for attendees to affix to their conference badges.

The green dot gave the impression that my acquaintance was a peering coordinator, responsible for negotiating interconnection agreements with other autonomous systems. Once I pointed this out to him, the reason for the questions he was being asked became clear. Peering coordinators evaluate other autonomous systems as candidates for interconnection by factors such as their size, geographic reach, and perceived reliability. Questions about the number of customers, and the history of acquisitions, were merely part of normal conversation amongst peering coordinators encountering one another for the first time, to establish an initial sense of these measures.

¹PGP (Pretty Good Privacy) provides cryptographic mechanisms for data encryption and authentication. Individuals who sign one another’s PGP keys establish mutual trust, which is used to build a broader web of trust. Key signing typically happens in person, so that identity may be reliably established.

²Sometimes detailed in presentations to newcomers at NANOG, such as <http://www.nanog.org/meetings/nanog52/presentations/Sunday/Lad-newcomers-intro-tutorial-nanog52.pdf>, last retrieved Feb 20, 2014. I borrowed the title for this section from slide 22 of that presentation, “Seeing Dots”.

My acquaintance quickly removed the green dot from his badge; he had affixed it under the impression that it was merely for decoration. As I reflected on this occurrence later, I realized that I had become an insider of sorts, familiar with the symbolic value of markers at NANOG meetings, and helping a newcomer to understand their meaning. I had many more conversations with my new acquaintance, at NANOG56 and subsequent NANOG meetings, during which he wanted to glean my perspective on NANOG and the personalities and relationships which inhabit it, just as much as I wanted to glean his.

This chapter is about the symbols, practices, relationships and technologies surfaced by this story, together forming a distributed system of governance which orders the topology of the Internet's inter-domain routing system. I examine these various features through the ways in which they are mutually constitutive with the network administration community, which gathers through NANOG meetings and the NANOG email list. In doing so, I explore the process of becoming a network administrator, through the eyes of my respondents, and through my own journey of enculturation into the world of NANOG.

The NANOG email list is one of the most important resources for network administrators around the world. All of my interviewees - whether in North America or South Asia - were familiar with it, even if they had never posted a message to it. Similarly, NANOG meetings are viewed as a critical location at which to meet with representatives of entities which are significant to the North American and global Internet, whether they be ISPs, data centers, Internet exchange points, non-profits, or large Internet companies such as Google, Netflix and LinkedIn. Interactions over the NANOG email list and at NANOG meetings form an important space through which coordination and collaboration occur amongst network administrators to order the Internet's inter-domain routing system.

I use the term "distributed system of governance" quite consciously, to bring "distributed" notions of coordination and collaboration into conversation with "centralized" notions of governance, which I often found pitted against one another in my interviews and observations. Following the discussion of trust in Chapter 2, I think of distributed governance as being constituted of a web of interpersonal and generalized trust relationships, supported by centralized institutional anchors (such as those discussed in the last chapter), evolving in relation to the risks and uncertainties of inter-domain routing described in Chapter 4. Distributed governance is a historically constructed set of social arrangements which are produced alongside the technological forms which they order, as I illustrated in Chapter 3. The web of trust relationships that I present here is complicated by economic and political interests, just as with my account of institutional anchors. Accordingly, the challenge I take on in this chapter is to engage with the ways in which these competing interests and obligations are worked out within the network administration community, as represented at NANOG.

In this chapter, I show how trust relations are produced and reproduced through the practice of network administration, and through meetings of the professional communities - such as NANOG - at which these practices are themselves produced and reproduced. The three concepts of trust, community, and practice are central to the account that I present here, each in themselves, and in their interrelation. I analyze the relationship between practice and community through the theoretical perspective of "legitimate peripheral participation"

established by Lave and Wenger (1991), which calls attention to the process of becoming a participant, as much as it does to the community of practice. I analyze the relationship between practice and trust in terms of the role that trust relations play in responding to the risks and uncertainties of inter-domain routing encountered in the everyday practice of network administration. I view community as the site at which trust relations are formed, especially in the face-to-face meetings of the network administration community at NANOG. I also analyze community in terms of its symbolic construction (Cohen 1999), to understand the norms and ideals which contribute to cohesiveness and membership in the network administration community. I re-engage with the problem of embeddedness of markets (Polanyi 2001; Granovetter 1985), by examining the role of the social relations of trust in the construction of markets for the interconnection of the networks that make up the Internet's inter-domain routing system.

The relationship between these theoretical concepts is by no means a static one. I open this chapter with a discussion of the history of NANOG, and the transformations it underwent as the Internet grew in the aftermath of the privatization of the NSFNET. I then show how NANOG meetings are critical sites for the production of trust relations, and for the development of shared understandings of practice. By examining the processes involved in becoming a participant at NANOG meetings, I explore the contribution of participation in the community of practice at NANOG towards the process of becoming a network administrator. Finally, I elaborate on the role of trust relations in the practice of inter-domain routing, for both the everyday practice of network operations, and the construction of markets for network interconnection.

6.2 Growing the Internet

When the NSFNET was privatized to form the Internet in April 1995, the NSFNET backbone network was the only autonomous system operating over wide geographical scales, requiring all other networks on the nascent Internet to interconnect through it. In short order, numerous other backbone networks were built, creating the beginnings of the complex topology of autonomous systems which constitute the Internet's inter-domain routing system. The transition from a hierarchical topology with a single backbone network, to a complex graph of autonomous systems, required substantial changes in technological form, shifts in the nature and density of relationships amongst network administrators, the production of personnel with new kinds of expertise, and the creation and evolution of institutional forms (detailed in Chapter 5). This is not simply a story of expanding scale - involving more people and more autonomous systems - but also one which takes into account qualitative shifts of changing balances of power, as economic and political interests began to vie against the technical interests which were the prime drivers of decision-making on the NSFNET.

In this section, I examine these changes, picking up where I left off in Chapter 3, to tell the modern history of the social and technical arrangements of the inter-domain routing system. I will show how NANOG transitioned from its original role as an enabler of coordination

and collaboration for the NSFNET-to-Internet transition, evolving substantially to become a key professional organization for those involved in the operation of the Internet. Through this process, NANOG became a critical site through which trust relations were constructed amongst network administrators, and at which understandings of practices were worked out, contributing towards the formation of the community of practice of network administrators.

6.2.1 Developing the Network Operations Community

Recall from Chapter 3 that NANOG was created as a NSF-funded activity to provide a space for coordination and collaboration amongst network administrators to help manage the transition from the NSFNET to the Internet. The first two NANOG meetings were held during June and October 1994 in Ann Arbor, Michigan. This was the home of NANOG's organizational sponsor, Merit.³ Subsequent NANOG meetings have been held in a variety of locations across North America.⁴

The earliest available archive of the NANOG email list from May 1994 illustrates the variety of conversations which took place amongst this emerging technical community.⁵ Some of these discussions were concerned with formal coordination activities. For instance, Merit staff asked for plans that autonomous systems had in place for the transition away from the NSFNET backbone to interconnection via the new Network Access Point (NAP) infrastructure.⁶ Merit staff were also involved in coordinating with autonomous systems to accept requests to add or modify entries in the Policy-based Routing Database (PRDb) from which inter-domain routing configuration for the NSFNET was constructed. Unlike the modern Internet, routing information was held at a single location, with updated routing configurations pushed out to routing infrastructure only once every few days.⁷

These conversations about coordination had the potential to spark sense-making discussions about what particular kinds of coordination entailed, and how coordination should be organized. For instance, a discussion emerged around one of the PRDb update notices, in which participants discussed how details of IP address block allocations should be maintained, balancing a centralized registry approach against a distributed approach leveraging rwhois servers, and detailing the different approaches taken across the North American and

³Founded as a non-profit corporation in 1966 to create network infrastructure amongst Michigan's three major universities. It was originally called MERIT - the Michigan Educational Research Information Triad - and played a key role in the NSFNET and the Internet's early history, managing the RADb amongst other endeavors. See <http://merit.edu/about/history/pdf/MeritHistory.pdf>, last retrieved Mar 1, 2014.

⁴For a complete list of NANOG meetings, see <http://nanog.org/meetings/archive>, last retrieved Mar 3, 2014.

⁵The earliest NANOG email list archives are available at https://www.nanog.org/maillinglist/mailarchives/old_archive/1994-05/threads.html, last retrieved Mar 3, 2014.

⁶NAPs were also a NSF-funded activity, providing common points at which autonomous systems could interconnect physical network infrastructure (presaging IXPs), to aid in the transition away from the NSFNET backbone.

⁷During May 1994, there were a total of nine notices of updates to the PRDb on the NANOG email list.

European regions.⁸

There were also posts detailing threats to the NSFNET infrastructure. These included notices and data from new tools to monitor route flap,⁹ notices of software vulnerabilities which needed to be addressed, and notices of failure (and subsequent recovery) of the ANS Network Operations Center (NOC) telephone lines.

Other posts dealt with the run-up to the first NANOG meeting. These included the agenda for the meeting, details for hotel reservations, and helpful posts with suggestions for travel arrangements. In addition, there were invitations to open houses and tours of facilities from the ANS NOC and private network service providers, along with invitations to social events, including whirlyball and splatball.¹⁰

The agenda of the first NANOG meeting reflected the primacy of coordination activities related to the inter-domain routing system: the entire first day of the meeting was devoted to presentations about NSFNET transition plans and the NAP infrastructure. Presentations on the second day dealt with the operation of route servers,¹¹ changes to the PRDb, the Internet Routing Registry system, and results from routing on networking testbeds.¹²

Even prior to the privatization of the NSFNET, conversations on the NANOG email list and activities at NANOG meetings covered a wide range of subjects, from details of network configurations, to discussions about policy and supporting infrastructure, to threats to network infrastructure, to discussions of formal and informal social occasions.

These conversations illustrate the emergence of a *community* of network administrators, patterned with thick relationships which can only in part be reduced to purely functional explanations of coordination and collaboration.

One of my interviewees, who is now a senior executive at a Tier 1 autonomous system, described his first NANOG meeting to me, held shortly before the privatization of the NSFNET:

The very first NANOG that I went to was NANOG3. It was in February of 1995. It was held at the National Center for Atmospheric Research here in Boulder, Colorado. It was in an auditorium there. The auditorium was only capable of seating maybe 200-250 people. It was only like one-third full. The thing that was unique about NANOG then was that almost every single person that was making the Internet work was in that room. If a bomb had gone off in that room, the Internet would have been set back a couple years while these other people came in to try and figure out how to pick up the pieces. It was that scary in terms of

⁸The rwhois approach for maintaining IP address block registration information continues to be supported by ARIN; see chapter 5 for details.

⁹See chapter 4 for an explanation of route flap.

¹⁰Whirlyball is a game similar to lacrosse, with players seated in fairground bumper cars. Splatball is a war game in which participants shoot one another with paint ball guns.

¹¹Servers which help manage BGP routing information at IXPs, which were originally put into practice at the NAPs.

¹²For the agenda of the first NANOG meeting, see https://www.nanog.org/maillinglist/mailarchives/old_archive/1994-05/msg00037.html, last retrieved Mar 3, 2014.

the - but also kind of exhilarating just to be in the presence of that much brain power. [I40:11]

This short description paints a vivid picture of the earliest NANOG meetings. They were small, with at most a hundred people in attendance.¹³ My interviewee drew a line around this group - the people making the Internet work - as opposed to those involved in setting protocol standards or managing resources. These are not mutually exclusive groups, of course; the early NANOG meetings featured presentations and representation from InterNIC, which was responsible for managing IP address block allocations at the time. There is, however, a distinction drawn to surface the group of people who hold knowledge of their autonomous systems' internal configurations and external interfaces to other autonomous systems. If this group were to be lost, my interviewee suggested, then the Internet itself might be lost, as it would take a while for others to acquire the situational knowledge and relationships of this group.

The excitement this interviewee felt in encountering this group of people comes across clearly. Many of my other interviewees reported similar feelings when they first found NANOG, at the mail list, and later at meetings, commenting on the kinship they felt with the NANOG community. Comparing NANOG3 and other early NANOG meetings to modern NANOG meetings, my interviewee continued:

The presentations were a lot more edgy. The conversations were a lot edgier, even to the point of being hostile. It was not uncommon for people to stand up and flame other people that were at the mics. In fact, even at those early ones they didn't even have mics. There was no PA system. It was just a guy in the front talking to an audience with no mic. There was no food back then either. There were maybe some cookies on breaks. I don't really recall. [I40:11]

There are two ways to interpret the edginess of conversations referred to in this statement. First, it could be that participation during presentations at NANOG called for a degree of aggressiveness, and a certain thickness of hide. Second, this could be viewed as a means to open conversation, albeit limited to those willing to engage in this kind of verbal sparring. As I will discuss later, this is not the only model of discussion at NANOG today, but it does remain an identifiable style of engagement, especially amongst those who have been part of NANOG for a long time.

The earliest NANOG meetings were short in resources compared to modern meetings. As my interviewee notes, they lacked basic infrastructure - such as microphones - and amenities. This changed as more people became involved in the Internet industry and attendance at NANOG increased. The earliest NANOG meeting for which attendance records are available, NANOG8, held in October 1996 in Ann Arbor, Michigan, had 156 attendees. The next

¹³In comparison, IETF meetings in this period almost always had 1000 or more attendees. To put this in perspective, it must be noted that the IETF is a global standards body, compared to NANOG, which is a regional professional association. For IETF attendance numbers, see <http://www.ietf.org/meeting/past.html>, last retrieved Mar 4, 2014.

NANOG meeting, NANOG9, held in February 1997 in San Francisco, California, had 502 attendees, a number which is representative of the attendance numbers at all subsequent NANOG meetings.¹⁴ My interviewee commented on the difficulties associated with this sudden growth:

Over the next few years it changed a lot because there were so many new people coming into the industry that had so little clue as to how things worked and how to get things done. A lot of the old timers that had been doing it for a while got kind of frustrated with that. They wanted the Internet to get commercialized quicker. They wanted things to kind of normalize. There's no way to do that unless you go through those growing pains, right? [I40:11-12]

Another interviewee echoed these concerns, drawing a clearer boundary between those involved in early NANOG meetings, and those who entered the industry later:

In the early group, people were very tightly knit, personally you have good relationships, so you don't mind sharing, but with other people who are not playing the game already, people may just say that's your problem, I'm not going to help you, or just get too busy to help pass really good information. [I5:14]

As the Internet industry grew rapidly in the late 1990s, a concomitant increase in the demand for operational personnel to manage Internet infrastructure placed stresses on the emerging operational community at NANOG. My interviewees' comments suggest that the growth in attendance at NANOG was not simply of more people with similar skills, but rather of people who needed to develop the skills and relationships that they needed to operate their autonomous systems. A period of socialization of new network administrators followed, which caused some friction between established community members and new entrants. This also resulted in a change in the programming at NANOG meetings. Earlier meetings had been primarily concerned with coordination activities, and documentation of operational and research experiences, featuring "State of the Internet" sessions with updates from NAP operators and the RADb, information about ISP network deployments, and research into network operations.¹⁵ As attendance increased, training became a regular feature of NANOG meetings. This began with a BGP tutorial introduced at NANOG9, which was described as follows:

Designed for new network operators and ISPs, the tutorial will provide information on practical BGP applications, hints, tips, tricks, and current implementations in today's ISP backbones. "BGP 101" won't be an-depth technical

¹⁴The attendee lists for NANOG8 and NANOG9 are available at <http://www.nanog.org/meetings/nanog8/registrants> and <http://www.nanog.org/meetings/nanog9/registrants>, last retrieved Mar 4, 2014. I discuss overall attendance numbers at NANOG in the sections that follow.

¹⁵This pattern holds for all NANOG meetings upto NANOG8. Details of past NANOG meetings are available at <http://www.nanog.org/meetings/archive>, last retrieved Mar 4, 2014.

presentation, so if you're already familiar with BGP, you may not wish to attend.¹⁶

Note how the description of the tutorial illustrates an understanding of varying levels of expertise amongst those attending NANOG, suggesting that those who are already familiar with BGP may not find it of use. Tutorials have remained a feature of NANOG meetings since then, expanding to cover a wide variety of subjects, and taking up to half a day of NANOG meetings.¹⁷ They have generally remained focused on *practical* information required for network operations, rather than purely theoretical knowledge of how to configure a particular vendor's networking equipment. Regular NANOG attendees have told me that the tutorials were well attended in those early years [I5:14], although they declined in utility as a core of NANOG attendees developed the requisite technical skills. This reflects a certain level of competence amongst personnel involved in network operations at large autonomous systems, and other significant entities (such as large content providers), involved in the operation of the inter-domain routing system.

6.2.2 Operating Unreliable Infrastructure

Failures of infrastructure were not uncommon in the early Internet. Stories of these failures continue to be circulated amongst the NANOG community, whether they be large publicly visible failures (such as the AS7007 incident detailed in chapter 4), or individual experiences of failures. These stories play an important role in the establishment of practice, as they provide a shared means through which to establish the boundaries of what is acceptable in network operations. These stories also provide the symbolic capital which allow individuals to establish their membership within the NANOG community, both in the telling of stories, and in the recognition of the import of stories.

For those involved with network operations in the early Internet, there is a certain relish to telling these stories. The failures they had to deal with were, no doubt, stressful and difficult at the time. However, in retrospect, they become tales of sense-making, detailing the formation of understandings in technical expertise, and the social relationships for coordination and collaboration required to resolve the failures which were part and parcel of the early Internet. These were as much failures of technology as they were failures due to a lack of expertise for the operation of technology.

For instance, a presenter at NANOG49 told the story of an ISP in Texas at which network administrators didn't realize that their IP address allocations were limited in size. They began numbering their network with IP addresses from the beginning of the block which was assigned to them, and kept going, expanding from a $\backslash 24$ to a $\backslash 23$ and then a $\backslash 22$. This went on for seven years, the presenter related - to much laughter from the audience - before they noticed that they couldn't get to some other locations on the Internet, since their prefixes were overlapping with those of the locations they were trying to get to [F-NANOG49:52].

¹⁶From <http://www.nanog.org/meetings/nanog9/bgptutorial>, last retrieved Mar 4, 2014.

¹⁷NANOG meetings originally lasted just two days, and were later extended to three days.

Stories of failure are as much about a lack of expertise - spanning internal operations and understandings of global conventions - as they are about technological problems. From this perspective, all failures must be viewed in varying combinations of social and technological elements, never purely as one or the other. Sketching his experiences with network operations on the early Internet, one of my interviewees raised several distinct concerns, which to him are all related in a continuum of experience:

There were relatively few customers on the Internet. They were pretty tolerant of interruptions to service at the time. Things were broken all over the place. You have to have a pretty flexible, agile environment in order to just get by. Yeah, it was pretty loose. It's why the engineers of the day got the titles known as cowboys because we were constantly shooting from the hip, doing exactly what I just said. Cisco said I've got a bug I need fixed. Cisco said the software will fix it. There was no lab to run that software in to do regression testing against previous bugs. There were no people available to do that. I was one guy trying to run a national network. I didn't have anybody else trained at that point in time yet that could do any of that intellectual work. The only option was to download the code and try it out. That was extremely common back in those days, back in the 1994, 1995 era. Later past that as well, I would say even into '96, '97, '98 that was still pretty common for people to run networks that loose. [I40:5]

BGP version 4 was only standardized in March 1995 (Rekhter and Li 1995), the month before the transition from the NSFNET to the Internet. Network equipment vendors - amongst which Cisco Systems was dominant - were building routers and developing router software to the BGP version 4 standard as the NSFNET to Internet transition was happening, and rolling out updates to software as they learned from challenges faced by network operators using their equipment. For their part, network administrators had to place their trust in the efficacy of Cisco's software releases. They had neither the facilities, nor the personnel, to test new software releases before deploying them to their networks. Network administrators were like "cowboys . . . constantly shooting from the hip", improvising solutions to immediate concerns, monitoring their networks as they rolled out changes to software, hardware, and internal and external topological arrangements.

Much of this was possible because the Internet was not then the critical infrastructure that it is today. While there was substantial uncertainty associated with network operations, the risks of failure were not as significant as they are now, if only because those beginning to use the Internet were more tolerant of failure. Another interviewee remembered this period:

. . . when BGP was young, the Internet was young too, and it wasn't considered a business-critical application. You know, there were days when the Internet was down the whole day. I mean the entire Internet, and it was not - I mean, yeah, people were upset, but it wasn't one of those things like you've just destroyed our livelihood like it is today. [I16:10]

Network administrators involved with operating ISPs on the early Internet had to deal with substantial uncertainties, in relation to the reliability of technology, and in relation to the ongoing development of skills for network operations, and the lack of personnel with these skills. The process of dealing with these uncertainties was a productive one, as network administrators - individually and as a group - came to shared understandings of their identity, and of the balance of concerns to take into account in the operation of their networks.

6.2.3 Circulating Relationships

The process of learning through failure contributed to the formation of a community of practice at NANOG meetings, and through the NANOG email list. While these were important sites at which learning was worked out and disseminated, there was a circulation of personnel amongst autonomous systems which aided the distribution of skills and knowledge across those involved in network operations on the early Internet. An interviewee involved with network operations in that period told me how personnel with the abilities to manage large autonomous systems were in short supply:

I'm just talking about the US, maybe 10 years ago, ...so there are a lot of little ISPs coming up, you know, there are just a certain number of people who are really experienced with BGP, and they cannot hire these people, either they cannot afford it, or these people are working with ISPs already, but they're not enough. [I5:14]

The resulting demand for their abilities led to a high level of labor mobility for these personnel. As they moved quickly from one organization to another, they passed on their knowledge and skills to those they encountered, establishing new relationships, and also carrying their relationships from NANOG and other organizations with them. An interviewee commented on this circulation of personnel, and how the resulting web of relationships helped in the resolution of issues in network operations, cutting across corporate boundaries:

... it is fortunate that, in the past decade, there has been a tremendous amount of movement of people between different network organizations. I mean, there was a time, not long ago, when if you found anybody who had been with a given network organization more than a year, that was a weird fluke, because everybody changed jobs every year so they'd get a nice bump in their salary. It was part of normal operations. Now that's sort of settled down, but for a while that was sort of common. So everybody knew everybody, from having worked with them once, pretty much, and so they had lots of personal contacts. If they couldn't find the information online some place, they'd probably say, well, I worked at this company, he worked at this company, and Jim still works at this company, I'm going to give Jim a call, and see if he still has Joe's current number. And then he calls Joe and says, Joe, you need to deal with this crap, you're really killing us! [I17:15-16]

Given the unreliable environment within which they were operating, network administrators had to be able to coordinate directly with one another to resolve issues, even when they had different employers. These cross-corporate relationships were especially important to addressing issues in the inter-domain routing system, in which failures at one autonomous system have the potential to affect other autonomous systems. These relationships cannot be reduced to purely functional explanations; they are carriers of more than simply coordination activity, as they form thicker bonds of friendship and trust, and help construct a larger sense of community, contributing to a more concrete community of practice. Accordingly, these relationships were used not just for responding to issues when they arose, but also to short-circuit corporate bureaucracies, and get mundane administrative tasks executed more quickly. A network administrator who was involved with the NSFNET, and worked with various ISPs after the NSFNET was privatized, told me:

The good thing about this Internet environment is, because in the early days, all of us, we all had a very good relationship with each other, I mean at that time the community was so small. So we went to IETF every 3 or 4 months, so we knew each other really well. Even though later we all worked for different ISPs and became competitors, we still had a good relationship. So it's very easy to communicate, you know, if people do filtering, you can call that person immediately, "I just need to be added to your filter, can you do it now?", and usually it can be done, right. [I5:10]

These relationships are complicated by economic concerns. How do network administrators reconcile their relationships with one another, and their affiliation to their community, with the economic interests of their employer? I will deal with this question in more detail in later sections, but it is worth examining here the ways in which the process of working out these interests affects the formation of practice.

The entity formed to manage the NSFNET backbone, Advanced Network Services (ANS), wielded substantial power within the early Internet, since almost all Internet traffic had to flow through ANS' network. Most of the ANS staff were drawn from those running the NSFNET backbone, carrying the practices that they had established in their operation of the NSFNET backbone over to the operation of the ANS network. Thanks to their central topological position, ANS could force requirements on anyone who interconnected with them, such as requiring the entry of routing data into RADb.¹⁸ As one of my interviewees, who was then (and continues to be now) involved with an autonomous system for a research institution, commented:

...in those days, one of the largest network providers, ANS, did their configs based on the RADb, so if you didn't put your data in the RADB, you didn't get routed by ANS, which was unacceptable. [I17:4]

¹⁸Recall from Chapter 3 that RADb - the Routing Assets Database - was the NSF-funded activity to maintain a database of routing information for all autonomous systems on the Internet.

Over time, other autonomous systems developed backbone networks of their own, and ANS personnel circulated to some of these other organizations, carrying their practices with them. As more backbone networks developed, ANS lost its central topological position; in consequence, its mandates no longer held global sway. Other autonomous systems could choose to maintain their routing configuration information in RADb, or not, as they saw fit. As I discussed in chapter 4, some autonomous systems chose not to use RADb, since public disclosure of their routing configuration in RADb could be used to infer their customer relationships, which they thought of as sensitive corporate information.¹⁹

Yet, other autonomous systems still continue to use RADb today, in spite of these concerns about disclosing customer information. This is sometimes a result of requirements placed upon them by their upstream network providers, who may require the entry of information into RADb to generate their routing configurations. This is also sometimes due to established practices within an organization, following a strong sense that using RADb to generate routing configuration is simply the right way to manage their routing configurations.

The development of practice is neither uniform nor linear. While maintaining more data publicly is attractive from a technical perspective - since it eases coordination activity - economic concerns may drive the same data to be maintained privately. The balance and history of these interests within an organization drive the form of practice adopted, as do the requirements which may be placed upon an organization topologically, by the policies of upstream network providers. NANOG and related venues provide the setting for presentations and debates, which work out the commonly acceptable variations of practice, and the conditions under which these variations are understood to occur.

6.2.4 Organizing NANOG

As I've already discussed, NANOG provides the institutional setting in which the social relationships, the community, and the practice of network administration are worked out. In this section, I explore the work that goes into organizing NANOG, and stabilizing it as a legitimate venue for the profession of network administration. The NANOG mission statement provides a starting point from which to understand the principles of organization at NANOG:

The purpose of NANOG is to provide forums in the North American region for education and the sharing of knowledge for the Internet operations community. NANOG is a venue in which technical matters pertaining to network operations and network technology deployment in Internet providers may be discussed among experts. Such discussions have in the past focused on, but are certainly not limited to, experiences with new protocols and backbone technologies, implications of routing policies on the Internet as a whole, measurement techniques and measurements of Internet health and performance, areas in which inter-provider

¹⁹This tension of private and public maintenance of customer information is strongly related to the similar issues which arose during the policy discussion at ARIN, related in chapter 5.

cooperation can be mutually beneficial (such as NOC coordination or security incident response), and maintaining a competitive and level business environment. NANOG serves as a bridge between the technical staff of leading Internet providers close to network operations, technical communities such as standards bodies, and the academic community. NANOG has consistently worked to maintain a high level of technical content in conferences and all related activities. In striving to achieve these goals, all tutorials and presentations, including BOF presentations, are reviewed in advance and are limited to those entirely of a general technical nature, explicitly prohibiting material that relates to any specific product or service offerings. For similar reasons, equipment exhibits are limited to specified special events at each conference.²⁰

Unlike the institutions of governance discussed in chapter 5, NANOG has no well-defined outcomes. It provides “forums”, and functions as a “venue”, for education, and “mutually beneficial” coordination and collaboration activity amongst the Internet operations community. In doing so, it also acts as a bridge between the Internet operations community, standards bodies (principally the IETF) and researchers. Although not mentioned in this mission statement, the bridging work at NANOG also brings in policy discussions from ARIN, ICANN, ISOC and the US government. The mission statement notably frames activity at NANOG as being important to “a competitive and level business environment.” This principle is enacted within NANOG through an explicit focus on presentations and sessions which are “entirely of a general technical nature”, prohibiting vendor-specific material to ensure that NANOG does not become a trade show. I will discuss each of these points in detail in the sections which follow.

NANOG continued to evolve alongside the Internet, sometimes reflecting shifts in the organization of the Internet (as with the addition of tutorials), and sometimes reflecting changing priorities within the NANOG community.²¹ For many years, Merit program management staff controlled all aspects of NANOG, from operating the website and mail list, to selecting venues, to deciding on the program at NANOG meetings. After a conflict with the program management staff, members of the NANOG community pushed for the establishment of a NANOG charter in 2005 devolving greater control of NANOG activities to the NANOG community. Merit continued to provide administrative and technical support, but control over the content and organization of meetings was transferred to a steering committee elected by NANOG attendees. The steering committee created additional committees to manage specific aspects of NANOG activities, including a program committee to manage the program at NANOG meetings, and communications committee to administer the NANOG

²⁰From the mission statement in the NANOG bylaws, available at <http://nanog.org/sites/default/files/sites/default/files/NANOG-Bylaws-%20October2013.pdf>, last retrieved Mar 7, 2014.

²¹For the discussion that follows, I draw from messages on the nanog-futures mail list, which is devoted to discussing the future of NANOG as an organization. See <http://mailman.nanog.org/pipermail/nanog-futures/2010-April/000742.html> and <http://mailman.nanog.org/pipermail/nanog-futures/2010-April/000748.html>, last retrieved Mar 8, 2014.

mail list and other forms of communication amongst the NANOG community. There were suggestions at the time that the NANOG community should take over all aspects of NANOG organization, but it was considered that the NANOG community lacked the infrastructure to take on these activities.

Concern about Merit's decisions regarding conference scheduling and venue selection, and the replacement of Merit administrative staff who were well-liked by the NANOG community, pushed forward discussions about shifting control for NANOG activities entirely to the NANOG community. The decision to consider separation from Merit was initially taken by the steering committee and an advisory group including members of other NANOG committees. Questions about this decision were raised on the nanog-futures mail list, as community members voiced concerns about the transparency of the process, but soon began to work out the details of a viable organization. A formal consultation was held between the steering committee and the NANOG community at NANOG49 in June 2010 to discuss the transition away from Merit. Eventually, a Delaware non-profit organization, "newNOG", was created as the vehicle for the new NANOG organization, supported by an initial loan of USD 250,000 from ARIN.²² NewNOG took over control of NANOG affairs from Merit in 2011.²³ NANOG as an organization is now financially entirely self-supporting: for 2013, it had expenses of approximately USD 1.5 million, against revenue of approximately USD 1.75 million.²⁴

The organization of NANOG today is governed by a board of six elected members, plus the executive director, who is employed to manage the day-to-day affairs of NANOG. The elected members are voted to their seats by the NANOG community, as are members of the various other committees, which include the communications committee, the program committee, and the development committee (responsible for fundraising activities). All board and committee members are volunteers, serving without compensation. The ability to vote in, and stand for, elections was originally determined by attendance at NANOG meetings. With the transition to self-governance, this is now determined by a paid membership, distinct from meeting attendance fees, although paid membership is not required to attend meetings. As with the IETF, membership is for individuals, not organizations, a distinction which was drawn during conversations about the new membership model.

The push for self-governance of NANOG's organizational activities must be understood in terms of the strength of the sense of community amongst those who attend NANOG. The NANOG mail list and NANOG meetings are the primary means through which NANOG is realized as a community, and more broadly, part of the means through which the profession of network administration is defined. The move by NANOG community members to take control of these means of existence helped further define commitments to their own notions of

²²The president of ARIN has stated that ARIN has an interest in viable network operators' organization in the North American region, and this was why they supported the NANOG transition to self-governance. The ARIN loan was repaid by the time of NANOG56 in October 2012 [F-NANOG56-ARIN30:50].

²³See the joint Merit-NANOG press release, available at http://www.merit.edu/news/newsarchive/article.php?article=20110201_nanog, last retrieved Mar 8, 2014.

²⁴From the NANOG 2013 budget, available at <http://nanog.org/sites/default/files/2013%20Budget%20website.pdf>, last retrieved Mar 8, 2014.

community, just as it furthered the conditions under which the NANOG mission, as written, could be acted upon by the NANOG community.

The community formed at NANOG came to provide an anchor to mitigate the risks and uncertainties of network operations, whether for technical issues requiring coordination and collaboration, or for social issues, such as managing relationships in the face of high labor mobility, or developing technical skills. In doing so, NANOG became a critical site for the production and reproduction of the skills, understandings, and relationships involved in the internal operations of autonomous systems, and the operation of the inter-domain routing system.

6.3 Going to NANOG Meetings

While there are a variety of formal presentations and sessions on the agenda, NANOG meetings are to a large degree a space for social networking, to construct and maintain the relationships that make NANOG a community. The social interactions at NANOG meetings - whether formal or informal - help maintain and develop the culture and practice of network operations, balancing technical and economic interests. Political concerns - such as the WCIT negotiations - are also brought into awareness, often through formal presentations. In the last section, I dealt with the historical evolution of NANOG as a community and as an organization. I continue to investigate these themes in this section, through an examination of the social world of NANOG meetings. As I will show, the production of the virtual space of the Internet is paradoxically reliant on face-to-face interactions in the shared physical space of NANOG meetings.

6.3.1 At NANOG

The first impressions I've had of every NANOG meeting I attended are of people socializing in the large lobby which provides the entryway to the main conference hall. Chairs and tables, and couches, are scattered throughout the lobby, which in some instances also features a hotel bar, creating an informal atmosphere for people to gather easily. Sponsored snacks, and some meals, are also made available in the lobby area. There are always small groups of people scattered across this space throughout the meeting, whether or not a session is in progress. NANOG meetings are as much about negotiating business deals, and catching up with old friends, as they are about the main conference sessions. These informal interactions also spread outside the main NANOG venue, with many people setting up appointments to go out to drinks or meals together, or hosting room parties.

ARIN maintains a help desk in the lobby area, to provide face-to-face resolution of any issues that attendees may have with requests to ARIN for resources. NANOG attendees, after all, are representative of the larger organizations which are ARIN members.

The dress code reflects the informality of NANOG meetings, with most attendees wearing t-shirts with shorts or jeans. Every meeting offers a sponsored meeting t-shirt as part of

the registration package. T-shirts can be a marker of seniority and membership, as some attendees wear t-shirts from prior NANOG meetings to show their long affiliation with the community. All attendees wear their badges throughout the meeting, providing attendees the opportunity to accost one another by name (in case they know one another only by email), or by organizational affiliation (to discuss technical or business questions). Markers on badges provide a means through which attendees may assess the kinds of conversations they may have: colored dots indicate particular job functions (such as peering or security), and badge color, or ribbons on badges, indicate position within the organization of NANOG, marking board and committee members, NANOG staff, and newcomers.²⁵

The main conference hall is arranged like that at ARIN and ICANN. A raised stage flanked by projector screens occupies the front of the hall, with rows of chairs facing towards the stage. Microphones are available in the aisles for attendees to engage in conversation with presenters, and sometimes with one another. Other sessions, held in different rooms, may be laid out differently; these are typically BoF sessions, which I discuss in detail later. Those who come up to speak at microphones are expected to identify themselves by name and affiliation.

The main conference sessions are mostly devoted technical presentations, covering issues of interest to network operators. These consist in large part of network operators presenting accounts of practical experience with particular technologies, ranging from inter-domain routing issues, to network design, to the challenges of deploying different kinds of physical infrastructure, such as fiber optic cables, or wireless links. Technical sessions are also delivered by computer science researchers seeking feedback and information for their research.

A number of other activities fill out the main conference agenda. Training sessions typically occupy a half-day of the conference. An update on the state of Internet resources is offered by ARIN, as is an ARIN public policy consultation on the occasions that ARIN is not held back-to-back with NANOG. Time is also set aside on the main conference agenda for a “community meeting” led by the NANOG board, during which the board reports on progress to the NANOG community. During the community meetings I attended, the organization of NANOG as an independent entity was a prominent subject of discussion, as were planned new activities, including the NANOG education program, a NANOG-on-the-road series of smaller NANOG meetings, and fellowship programs. As new positions open up on the board and other NANOG committees, the community meeting also serves as a space for attendees to pitch their candidacy to the NANOG community.

The process of entry into participation in the NANOG community has over time become formalized to a degree. A special breakfast or lunch event is provided for newcomers, to introduce them to the world of NANOG. This is organized in a separate room, with tables at which long-time attendees are seated alongside newcomers. Unlike ICANN, ARIN or the IETF, there is no formal process or policy to explain; the goal of the newcomers meeting is to help newcomers understand the social norms and relationships at NANOG. Newcomers

²⁵The welcome mail for NANOG57 details these markers, see <http://mailman.nanog.org/pipermail/attendee/2013-February/000128.html>, last retrieved Mar 15, 2014.



Figure 6.1: The view from the audience at NANOG56.

meetings sometimes (but not always) feature a formal presentation, in which aspects of NANOG are decoded for newcomers, such as the meaning of dots on badges.

Commercial activity is carefully regulated at NANOG, to ensure that all NANOG activities are in service of the development of a better Internet, rather than for the commercial gain of individual vendors or network operators. Recall that the NANOG mission statement explicitly prohibits vendor-specific material from NANOG conference presentations, and restricts vendor activity to specific venues at the conferences. That said, those who attend NANOG regularly, and especially those who serve in the organization of NANOG, are quite aware that NANOG requires sponsorship in order to survive, and needs to ensure that sponsorship opportunities are of interest to potential sponsors. NANOG presents a significant marketing opportunity for many networking equipment vendors, since NANOG attendees are typically senior technical personnel who are involved in purchasing decisions.

All sponsors are thanked during the conference opening and closing. The only sponsors who are provided with time during the main conference sessions are the sponsor acting as

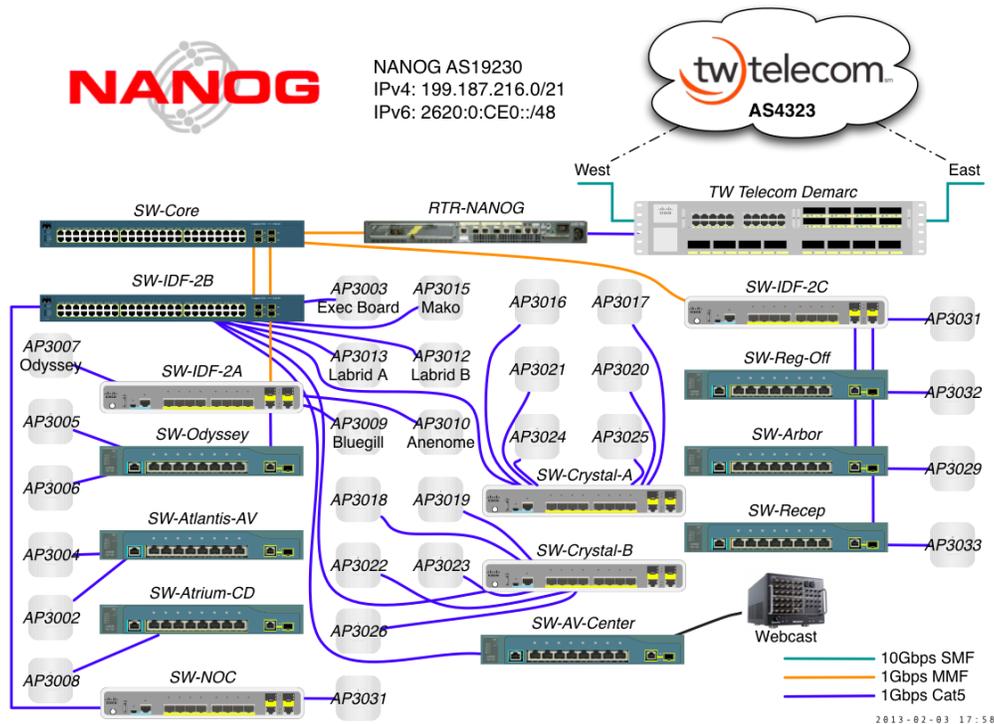


Figure 6.2: The network for NANOG57, provided by connectivity sponsor TW Telecom, showing wireless access points, routers, and redundant upstream connectivity, available at <https://www.nanog.org/meetings/nanog57/diagram>, last retrieved Mar 18, 2014.

conference host, and the connectivity sponsor. Both of these are offered this opportunity to talk about their service offerings to the NANOG audience, which typically include global and regional maps of their facilities, and an indication of their interconnectivity with other autonomous systems. The connectivity sponsor also provides a schematic of the network that they've set up for the NANOG conference, from the wireless routers in different conference rooms to the router aggregating these links and providing upstream connectivity through the sponsor's autonomous system. It is a point of pride to show off the capabilities of networks, such as IPv6 support, which is always made available on the NANOG conference network (see figure 6.2 for an example).²⁶ NANOG meetings always feature one or more large off-site social events, at least one of which is typically sponsored by the conference host. During the NANOG meetings I attended, these have been at venues which have ranged from nightclubs, to a steamboat (in New Orleans), to a section of Seaworld (in Orlando).

Vendors are provided space to advertise their wares at the “beer 'n gear” event. This event takes place in a separate room, in which vendors occupy booths around the edges of the room, with drinks and food made available for attendees. Like other parts of NANOG

²⁶NANOG has its own autonomous system number and IP address space which is deployed for the NANOG meeting network, as indicated in figure 6.2.

meetings, this is as much an opportunity for attendees to socialize as it is for them to interact with vendors.

Commercial activity is not only restricted to the formalized spaces and sessions which are provided to vendors and sponsors. NANOG is a gathering of personnel who are often involved in both technical and commercial activity in their organizations. Consequently, there is work that goes into the definition and construction of boundaries between the commercial interests of individual organizations, and the more general technical/economic discussions which are considered acceptable at NANOG. For instance, it is acceptable to talk about a generalized model for the total cost of operation of a carrier-grade NAT system,²⁷ but not to turn a presentation into a marketing spiel for consultancy services to apply such a model.

It can be tricky to navigate the boundary between acceptable and unacceptable commercial activity. Opinions about whether or not this boundary has been transgressed have generated contentious discussions within NANOG in the past, and sometimes have even led to a degree of ostracism against those who have engaged in what has been perceived as unacceptable commercial activity [F-NANOG56-ARIN30:218].²⁸ These issues are subjects of discussion on NANOG email lists as well, with participants reflecting on how the boundary between acceptable and unacceptable commercial activity sometimes becomes clear only once their peers call them out for transgressing this boundary.²⁹

This boundary is constantly drawn and worked out in practice at NANOG meetings. At NANOG56, for example, a NANOG committee member overseeing the training sessions thanked the trainers, but drew attention to the fact that vendors were involved in delivering the training, “this is something for which you may have to rely on a vendor, which is not always neutral” [F-NANOG56-ARIN30:100]. Similar concerns were raised at the NANOG58 community meeting, during which an attendee asked whether there were any concerns about NANOG’s independence, pointing out that 70% of the NANOG budget comes from vendors. Board members responded that they saw no cause for concern, since “sponsors have been kept out of the program”, commenting that “sponsors know what they’re getting into; no-one asks pay-for-play, no-one asks for a speaking spot.”

The formal structure of NANOG meetings covers a range of areas, providing plenty of space for attendees to socialize with one another, and find their way into the community at NANOG. The diversity of activities at NANOG mirrors the diversity of attendees’ positions and interests, and is maintained as a space for technical discussion and community through careful segregation of commercial interests.

²⁷Such as in this presentation at NANOG56: <https://www.nanog.org/meetings/abstract?id=2025>, last retrieved Mar 18, 2014.

²⁸I refrain from providing specifics of instances of transgression, since these are sensitive topics.

²⁹For instance, see this message on the nanog-futures mail list, available at <http://mailman.nanog.org/pipermail/nanog-futures/2010-July/000878.html>, last retrieved Mar 14, 2014.

6.3.2 Attendance Patterns

NANOG meetings occur three times a year at locations across North America. These shifts from one location to another are intended to maximize participation from smaller local ISPs who are typically unable to attend NANOG on a regular basis. NANOG attendees are mostly from medium to large organizations operating Internet infrastructure in the North American region, although there are attendees from other parts of the world as well, often from IXPs at which the organizations represented at NANOG interconnect. For all that the Internet is a global system, organizations from North America remain significant providers of global connectivity. At the NANOG meetings I attended, most attendees were from the USA, but twenty or more countries were represented in total.³⁰ Accordingly, while NANOG meetings are primarily devoted to working the professional community of network operators in North America, they also represent a site at which social and commercial relationships stretching to other parts of the world are developed, and at which understandings and practices of the profession of network operations are worked out, which may have global implications.

Like the institutions discussed in the last chapter, NANOG meetings are open for anybody to attend, but regular attendance requires commitments of time and money, which typically call for an organizational sponsor. Attendance fees for NANOG meetings currently stand at USD 600, to which must be added costs for lodging and travel.³¹ Drawing from available attendance records, attendance at NANOG meetings has varied between a low of 156 (at NANOG8, held in October 1996 in Ann Arbor, Michigan), and a high of 689 (at NANOG59, held in October 2013 in Phoenix, Arizona), with an average of 491 attendees across all meetings from NANOG8 to the most recent NANOG60.³²

These numbers tell only half the story. Commenting on the different social groups at NANOG, a frequent NANOG attendee told me:

The perception seems to be that NANOG's become routine. That the peering coordinators show up, they don't really bother going to the meetings. They set up the meetings with the peers they need to talk to while they're there. They go out and they party in the evenings and do a lot of relationship building. Yeah, you know, maybe they go to a presentation or two and learn something, but they're really there to meet with peers and to get things done in that regard. That's almost like a subculture of NANOG, if you will. There's a whole other aspect of NANOG which is some of the old timers contributing their time to give a talk on a certain subject and then a lot of newcomers coming in trying to learn

³⁰These numbers are published by NANOG staff and presented at the conference closing.

³¹NANOG has a discounted student registration fee of USD 100, and also offers two fellowships which cover all costs of attending a meeting. See <https://www.nanog.org/resources/fellowships>, last retrieved Mar 12, 2014.

³²All data is drawn from attendee lists linked to individual meetings, publicly available on the NANOG website from NANOG8 onwards. Attendee data is not available for NANOG10. I also use this data to generate the attendance graphs which follow.

about this new crazy world of how the Internet functions, as if it's some kind of hacker community. They're striving to be part of the inside club.

... [presentations] I've seen have shown that most of the newcomers are not coming back to a second NANOG. A third of the crowd is new every single time, which means that your core isn't really changing. You have a core of people that are actually involved in operations in the Internet that actually have businesses that pertain to running a network and they need a place like NANOG, much like a peering forum, to come and to execute in the function of their job and their role and to do what needs to be done for their network. That core, although people have come and gone into it as companies have hired new people and trained them up to become peering coordinators and the old peering coordinators leave to go and do something else or go to some other company. It's expanded over time as you would expect, but there's still this core of NANOG and then there's this periphery of NANOG around it. I don't recall the delineation between the core and the periphery being so distinct a decade ago as it is today. [I40:12]

This interviewee, and others, pointed to several interesting fractions of the NANOG community. First, peering coordinators, who are at NANOG to negotiate business agreements for interconnection. These are bilateral agreements, requiring one-on-one meetings and socializing outside of the main NANOG meeting. Second, "old timers", presenting talks on technical subjects of interest to the NANOG community. Third, newcomers who are trying to enter the "inside club". To these I add a fourth fraction, who occupy an interstitial position between newcomers and old timers, having successfully socialized into the NANOG community, but not yet achieved the seniority and reputation to be considered old timers.

I analyzed NANOG attendance records to see if these fractions could be observed from attendance patterns. Figure 6.3 and figure 6.4 show the changes over time in overall attendance at NANOG meetings, the numbers of new attendees, and the numbers of new attendees who only attend one NANOG meeting. As figure 6.4 indicates, there was a greater proportion of new attendees in the late 1990s than there is today, which reflects the growth of the early Internet, and the comments I referenced earlier regarding the sudden influx of newcomers to NANOG in that period.

Peering coordinators, old timers and the interstitial fraction are all part of the core attendance at NANOG meetings, which could be considered a proxy for the membership of the NANOG community. My interviewee asserts that this core is largely invariant. I analyzed NANOG attendance figures to compute the core of NANOG attendees, and changes in membership of this core, as illustrated in figure 6.5 and figure 6.6.³³ My analysis suggests that my interviewee's assertion is correct: the core of attendees is relatively stable, if the core is defined as individuals who attend four or more meetings (out of six) in a two-year period. I also consider a different threshold, for individuals who attend more than two, but less than four, meetings in a two-year period. For this lower threshold, greater variation

³³All computations are for a two-year period, with results plotted at the central point of the period.

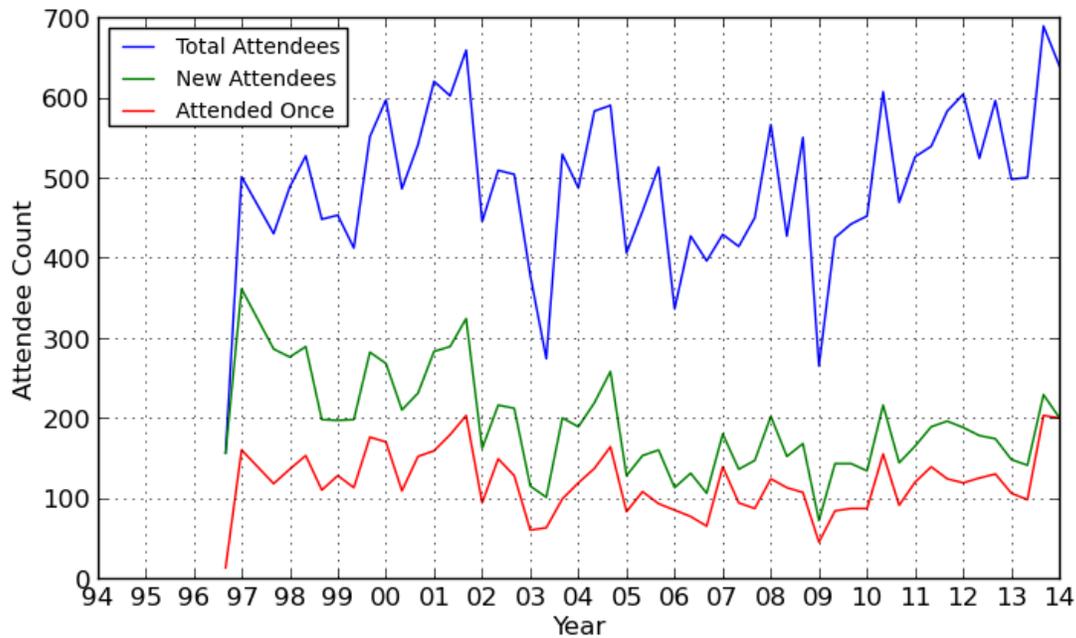


Figure 6.3: Comparing total meeting attendance against counts of new attendees, and counts of attendees who attend only one NANOG and never return.

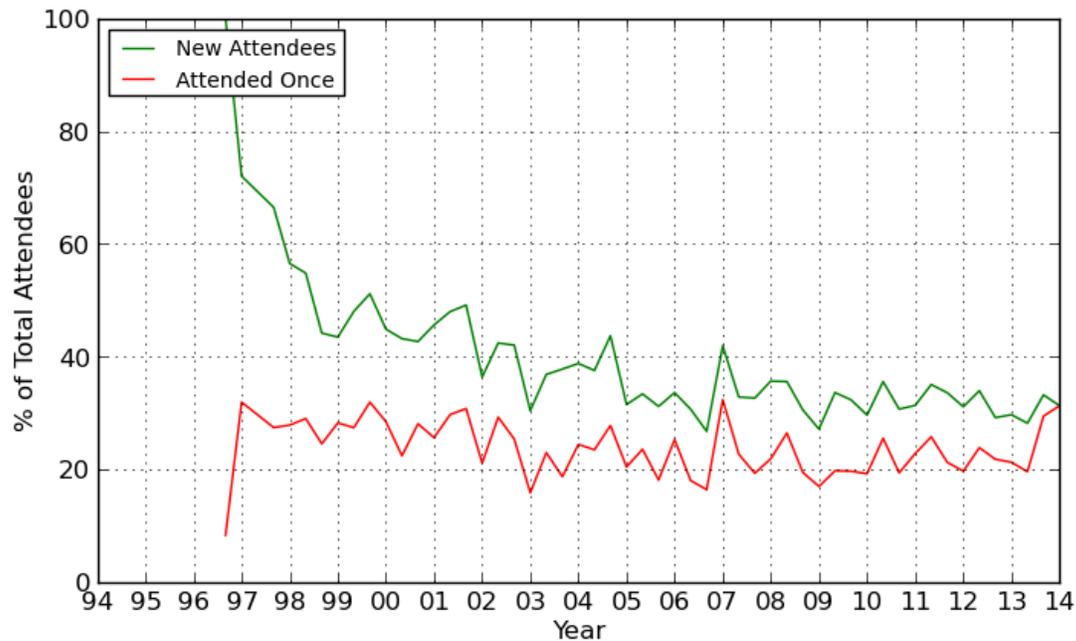


Figure 6.4: New attendees and attendees who attend only once, in proportion to total meeting attendance at NANOG meetings.

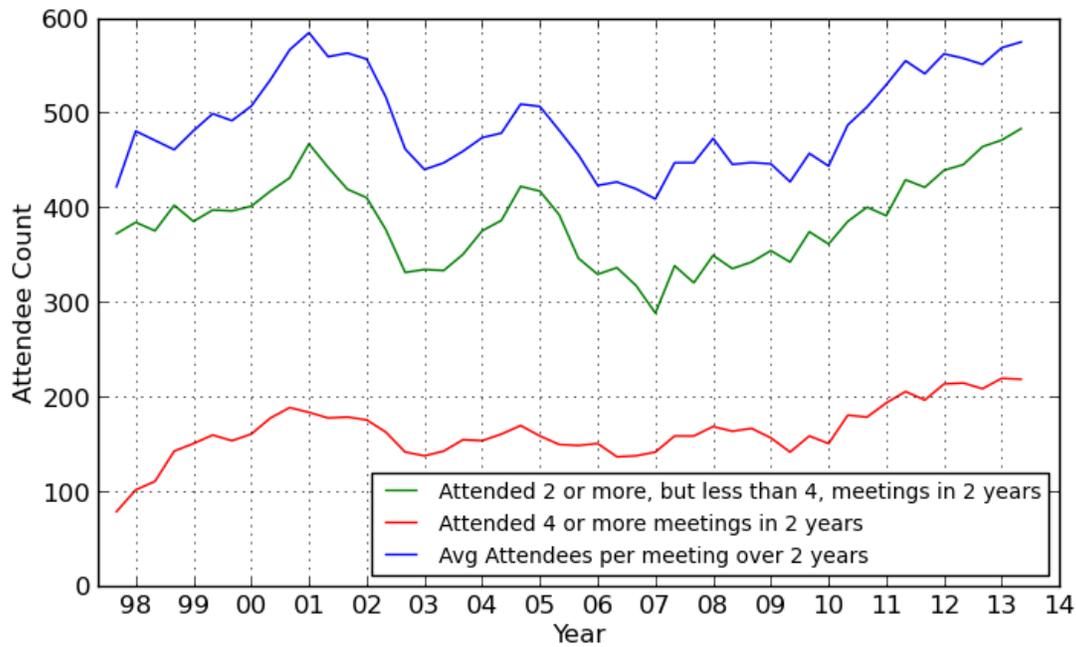


Figure 6.5: The core of NANOG attendees, computed as attendance at two or four meetings over a two-year period.

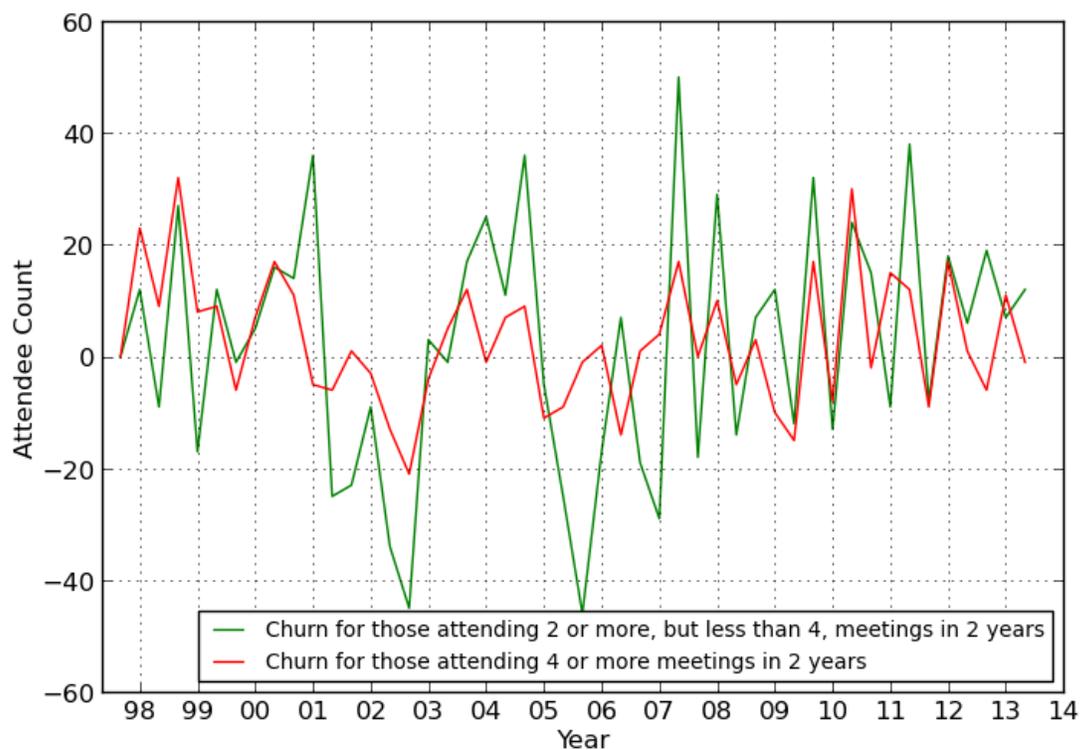


Figure 6.6: The churn in the NANOG attendee core, calculated as the number of new members of the core group minus the number of members who left the core group.

becomes visible in the size and membership of the core of NANOG attendees. In both cases, there is continued change - albeit at different magnitudes - in the individuals constituting the core attendee group, as illustrated by figure 6.6, which tracks the difference between the number of new entrants to the core attendee group, and the number of people who left the core attendee group.

These patterns are likely a consequence of funding considerations for attendance at NANOG meetings: small-to-medium organizations may send their personnel to only one meeting per year, while larger organizations may have the wherewithal to send their personnel to multiple meetings, and see the value in doing so. In addition, organizations of all sizes may send certain key personnel - such as peering coordinators or senior technical personnel - to NANOG meetings more frequently than other personnel. The consequences of funding are quite visible in figure 6.5, which shows a drop in overall attendance and core attendees from 2001, right after the Internet bubble burst. Independent consultants make a point of attending NANOG regularly, since this is where they find a market for their skills. Finally, there are some individuals who are part of the core group of repeat attendees, who come because they see NANOG as a gathering of their friends, of people with whom they feel a kinship, in sharing interests in similar technologies, and sharing dependencies on one another to get their work done. In interviews and conversations with those who are part of the core of NANOG attendees, I was repeatedly told - often with pride - how they had almost never missed a NANOG meeting since they started coming.

Apart from job function and position within the NANOG community, there are a set of considerations of identity - race, gender and age - which also shape the character of NANOG meetings, and are recognized concerns within the NANOG community. At NANOG58, I attended one of the few non-technical presentations I've seen at NANOG, presented by a senior member of the NANOG community and a gynecologist, on working from home. Amongst the points they made was that diversity of culture and experience makes for a more creative workplace, leading one of the presenters to gesture towards the audience and say, "In case people are confused, this is not what diversity looks like", eliciting laughter in response. He made this statement over a presentation slide with a footnote, "Look to the left. Look to the right. If it's another white male geek, you might be performing suboptimally." [F-NANOG58:407-426]³⁴

As the presenter pointed out, attendance at NANOG is overwhelmingly male and white. A female attendee at NANOG58 told me how another longtime female attendee was trying to organize all the women attending NANOG, to support one another. She joked that since there were so few of them, perhaps they should call themselves NetBFFs [NANOG58:101]. This is part of a long-run pattern: one of my interviewees related how a Merit program manager for NANOG used to count the number of women coming to NANOG meetings, and stopped once it dropped below a certain number.

NANOG is also aging. At NANOG56, during a presentation by a representative from

³⁴Presentation slides and video are available at <https://www.nanog.org/meetings/abstract?id=2132>, last retrieved Mar 15, 2014.

the US Department of Homeland Security (DHS) working on cybersecurity, the presenter asked all attendees younger than thirty to stand up. Out of an audience of almost six hundred people, about twenty stood up. The presenter noted that this was a better number than in many network security communities he engaged with, but showed how urgently younger people needed to be brought into the network operations community [F-NANOG56-ARIN30:117]. I discussed this pattern with several NANOG attendees afterwards, all of whom expressed concern over the lack of younger network operators at NANOG, but some of whom were also pragmatic about this pattern [F-NANOG56-ARIN30:193]. Many of those who are regular attendees began coming to NANOG when they were much younger, and learning how to start and operate networks in a time when the commercial organization of Internet infrastructure was still relatively informal. As the organizations operating Internet infrastructure have grown larger, and formalized their internal structures and processes, it has become harder to justify sending younger network operators to NANOG meetings.

There are multiple segmentations running across the attendance at NANOG meetings, which may broadly be divided into two types. The first is by job function and the relative position of an organization within the infrastructure of the Internet. This leads to segmentations in types of activity at NANOG, and in attendance patterns, which shape the nature of community at NANOG. The second is by identity, which illustrates the lack of diversity at NANOG. I shall be more concerned with the effects of the first of these types of segmentations than the second, but the second is worth noting to help understand the nature of NANOG meetings.

6.3.3 Participation

In the last chapter, I discussed participation in terms of well-defined outcomes and process for the institutions that I examined. Unlike these institutions, NANOG has no well-defined outcomes or process. As the NANOG mission statement indicates, it is a “venue”, providing “forums”, for education, coordination and collaboration. In what terms, then, can an attendee’s “participation” in NANOG be understood?

As I’ve already noted, NANOG is open to participation, with no prerequisites for registration at the conference, or subscription to the mail lists. However, as with the IETF, meaningful participation requires a recognition of expertise and position, in ways which are defined by the sociotechnical relations of interdependence required in the practice of network operations. For the purposes of analysis, I separate participation in conference sessions from participation in social gatherings, and then show how these two forms of participation are connected. I will show how participation may be understood in terms of individual reputation, job position and organizational affiliation, and the relative significance of an organization within the inter-domain routing system. From this perspective, participation may have different definitions: a peering coordinator will have a different account of participation than a NANOG committee member, and the account of participation from a small ISP will, no doubt, be quite different from both of these.

In the previous section, I discussed the segmentations at NANOG meetings, showing how there are multiple layers amongst those who attend NANOG, from attendees who come to almost every NANOG, to others who make it about once every year, to a not insignificant proportion of first-time attendees, many of whom may never attend a NANOG meeting again. The most peripheral kind of participation at NANOG is that of personnel from small ISPs, who may be attending NANOG for the first time, and may never return. Their experience is often one of attending training sessions and presentations, and possibly meeting personnel from their upstream network providers, and network equipment vendors. At the other extreme of participation lie personnel from medium to large Internet infrastructure providers. They have typically attended multiple NANOG meetings in the past, and have established relationships with other regular attendees. Their interests are in the maintenance of these relationships, so that they may leverage them for everyday network operations - being able to call someone in another organization directly to get things done - and for the negotiation of peering agreements. As I've discussed previously, these relationships are not purely functional in themselves, although they may be leveraged for functional outcomes. Through repeated interactions, evaluation of each other's technical expertise, and reliance on one another, attendees form thicker bonds of trust and friendship.

Although the gathering in a common physical space at a conference might encourage analysis in terms of physical engagement, participation at NANOG conference sessions should by no means be understood purely in terms of face-to-face interactions. It is rare to find an attendee without an open laptop during conference sessions, attending to email, or engaging in backchannel chat sessions with other attendees. Many attendees continue to work during the conference, actively monitoring the network infrastructure they are responsible for.³⁵ During a tutorial at NANOG58, the presenter asked the audience how many of them didn't have a pager or some other means of being contacted urgently. Of about a hundred people in the room, just three raised their hands [F-NANOG58:17]. Those attending NANOG who are involved with network operations are constantly on call, and they carry connections to their work with them no matter where they are at the conference.

The multiplicity of interactions that attendees may engage in through different spaces at the conference manifest visibly in the singular space of the main conference hall: attendees are engaged in individual work, conversations with one another, attention to the presentation in progress, and interaction with the presenter. These varieties of interaction are not mutually exclusive; each feeds into the other, sometimes resulting in behaviors which breach conventional boundaries - such as that between presenter and audience - but which are constitutive of the sense of relationships, reputation and community at NANOG. For instance, during a presentation by the president of the NANOG board at the NANOG56 community meeting, a senior and well-respected NANOG attendee wandered slowly up to the stage to speak to a board member. The president of the board was distracted for a moment - as were many in the audience - but she recovered quickly, joking, "He's wearing his invisibility

³⁵Attendees may even monitor the status of the network infrastructure provided for the NANOG meeting, using tools available through the NANOG website.

cloak” [F-NANOG56-ARIN30:72]. The acceptance of the attendee’s action, and the way in which it was brushed aside, can be viewed as a result of the regard in which he is held; after all, not everyone can have an “invisibility cloak”. However, it was the relatively informal air of the proceedings which made his action possible, since the board members are viewed as not being separate from the community, but rather as community members who have volunteered themselves to service for the community.

The separation between presenter and audience is similarly fluid throughout all NANOG sessions, as many presenters are well-known to regular attendees, with their relationships shaping the sometimes combative nature of their public interactions as well. Presentation at NANOG meetings, and participation in responses to presentations, are some of the means through which attendees can gain reputation within the NANOG community. One of my interviewees elaborated these themes, providing a unified account of trust relations, functional outcomes, and the management of individual reputation:

Certainly NANOG is one of the places where you can go and meet them face-to-face, where you can become a trusted individual, where you can get up and make a presentation about how cool your network is, what are the cool things you’re doing. It’s also a place where you can get up and talk about problems that you’re seeing in the network, to get other people together that are seeing the same problem, to generate more push towards vendors to solve the problem, or if it’s a process issue, to generate more push to get the processes changed, and ISPs work with each other. I see that as very important. . . . Other than that, I think NANOG is a place where people can go and show off how smart they are. [I14:10-11]

The importance of face-to-face interactions for the establishment of social relations at NANOG meetings was a recurring theme in many of my interviews. A long-time NANOG attendee related his experience of the formation and use of social relationships at NANOG meetings:

So we get along, we talk well. If you go to meetings on a regular basis, you get faces to go with names. . . you see names, you’ve got email addresses in your archive, and you keep some of these, I don’t keep every single message but I keep a lot of them, and I certainly have addresses for lots of people that I may need to contact on short notice. And it just makes the system work surprisingly well for a system that was never planned, it just happened, it just grew organically. Unlike many things that grew organically into this uncontrollable mass that is useless, it’s continued to be fairly effective. I think mostly because it stays fairly lean. [I17:16]

These face-to-face interactions occur across a range of venues outside the main conference sessions, sometimes by prior arrangement, sometimes simply as part of a group that’s gathering people to go out together, and sometimes spontaneously at various social events.

During one of the sponsored social events at NANOG58, a network administrator who I'd met at an earlier NANOG meeting introduced me to an acquaintance of his. They were meeting for the first time in person, although they had exchanged many emails in relation to setting up BGP peering sessions between their autonomous systems. Their discussion was a mix of technical concerns, and personal background, as they got to know one another socially, rather than simply as employees of organizations setting up peering sessions with one another [F-NANOG58:136]. Reflecting on the confidence he gained by the second NANOG meeting he attended, a young network operator told me:

I think for the first time I came to NANOG, like any other who was a first-time attendee, I was kind of an introvert to be asking questions and making connections, when you see the others with very different backgrounds, both in terms of age, professional experience, and the domains that they come from. So it was a little difficult for me to reach out to them and speak to them. I would also feel that I wouldn't have much to share or contribute to a conversation. . . . I would say when I had the first-time attendee sticker on my badge it was easier to talk to people. They would come to you and talk and see what you do and which company you work for, your role and stuff. . . . Maybe it's not just because of the badge or other things that might be affecting it, but this time it's less of others coming to me and talking to me, but it's more of me maybe going to them and speaking to them and getting introduced to someone. [I49:14-16]

During the second NANOG meeting he attended, he had been tasked with meeting personnel at specific other organizations, but was also generally more confident about approaching other attendees. Another interviewee related how these face-to-face interactions are of importance to building relationships, which may later be leveraged to ease inter-organizational collaboration:

It's an interesting opportunity to kind of put a face to a lot of people who I end up establishing relationships with, so I mean we'll end up establishing relationships because most of the communication that I have is usually done through email, mostly - seldom we'll get on the phone, if we're trying to establish a bigger. . . if we have a maintenance to kind of work out, but a lot of it's done through email, and there's good reason behind that. . . .this kind of conference gives us the opportunity to kind of say, "Hey, thanks for peering with us. It's great. Let's go grab a beer." That's kind of what we're looking at for coming here.

. . . It makes things a lot easier in the long run to have those conversations, so if for some reason you need to - there's something that you need, or internally maybe one of your - maybe you're having a capacity problem, and for some reason you need to upgrade a link, they might not be ready to upgrade a link, but if you - you can have that conversation. You can call them up, and say, "Hey. How's it going?" You've talked to them before. They understand. They know

you're not trying to pull one over on them. You're being honest with them, so usually that conversation goes a lot easier than if you haven't had the face-to-face conversation with them, so yeah. Ultimately it is a lot easier to kind of manage that relationship. [I19:11-12]

The informal conversations at NANOG meetings enable the sharing of knowledge which may ordinarily stay sealed within corporate boundaries. As a long-time NANOG attendee told me:

NANOG's always been a very open community. It's scientific in a way, right, where we all feel that it's much better to publish our failures so that everyone can learn from them. In turn, by everyone publishing their failures, then we learn from their failures as well. . . . We're willing to compete on who has the best price and who's network works the best as opposed to artificial political ways of trying to set up a competition. I'm not saying our companies necessarily follow that, but we as engineers definitely do. There's been - you hear all kinds of talk in hallways where people are talking about things that their companies would be very disappointed or may even violate their agreements with the companies, but they do it anyway because we're engineers. [I44:7]

Participation at NANOG meetings is as much defined through interactions in informal settings, as it is through interactions during the formal conference sessions. Both of these kinds of interactions contribute to the formation of social relationships amongst conference attendees, although with different effects. Interactions during formal conference sessions contribute to the formation of a relationship between a commenter and the presenter, but are also publicly visible to all present, generating well-articulated perspectives on the viewpoints presented, and individual evaluations of technical competence. Interactions in informal settings contribute to thicker interpersonal relationships, as attendees get to know one another socially, and not merely in the terms of their job functions.

My experience at the first NANOG meeting I attended, NANOG49, was in some ways similar to that of a first-time attendee from a small ISP. I knew almost no-one personally (I did come across a few of my interviewees), and I held little status, not even that of a customer, which a small ISP might hold. I was, however, a student, a researcher from UC Berkeley, a status which carries a valence of its own. Almost every NANOG meeting has some academic presence, typically computer science researchers making presentations to get the network operator perspective on their work. Academics are a minority at NANOG, but are certainly not alien to it; further, students are offered a substantial discount on the conference fee, marking students as a recognizable category amongst NANOG attendees. This familiarity with the category I belonged to, and my status, carried over to my interactions at NANOG meetings. For the most part, attendees were quite willing to engage with me in casual conversation, and take time out for formal interviews during the conference, or set up an appointment for an interview at a later time. This was as much a reflection of my status

as student/researcher, as it was of the relative openness of NANOG to casual conversation, and associated professional networking opportunities.

During the early phase of my fieldwork, many of my conversations and interviews were opportunistic, as I sought to engage with attendees I encountered around a shared breakfast table, those who were seated next to me during the main conference sessions, or attendees I met during the various social events that were part of the conference. My early engagements were somewhat naive, as I had to look up terms and stories that came up in my conversations. Although I have a technical background in computer science, many of the considerations involved in operating autonomous systems, and managing interconnections between autonomous systems, were quite new to me. For instance, I first learned about the RADb and the IRR system through interviews, and had to later spend time researching them. Similarly, as people told me the shared stories which are circulated amongst network operations personnel - such as the AS7007 full table leak, detailed in chapter 3 - I began to form my own frame for understanding and engaging in conversations at NANOG meetings. No matter how much technical training an attendee has, becoming a network administrator - especially at a medium-to-large autonomous system - requires additional socialization into the practice of network operations, and the world of NANOG.

As I attended more NANOG meetings, and encountered attendees I had met at previous meetings, I began to be invited out to lunches and dinners, or invited to join attendees I knew at their lunch table in the conference lobby area. These encounters gradually became self-reinforcing, as attendees I had met before introduced me to their friends, some of whom I had already spoken with. In such instances, both those doing the introducing, and those I was being introduced to, took a second look at me, reappraising me in light of the shared connection that the introduction surfaced. By the time of NANOG58, the last NANOG meeting I attended for my fieldwork, I was being greeted by acquaintances who said, laughing, “Oh, you’re still here?”, and encouraged me to bring the results of my research back to NANOG [F-NANOG58:434]. Other attendees that I’d met at previous NANOG meetings began to have conversations with me that were no longer restricted to discussions that were just about network operations, but rather about exchanging life stories. I learned where people came from, heard stories about marriages and children, and told similar stories about myself in turn.

This depth of interaction was, of course, limited to a relatively small set of attendees. I experienced a variety of responses in engagements with attendees at NANOG, problematizing the openness of interaction at NANOG. In addition, NANOG meetings often feature private parties, which I only heard about in passing. These variations in engagement may be understood in terms of the position I occupied relative to those I engaged with, and in terms of their organizational affiliations and associated economic interests.

During NANOG57, I introduced myself to the peering coordinator for a Tier 1 autonomous system, to ask her for an interview. She responded that she was unable to give interviews without first clearing interview questions with her legal team, holding up her coffee mug to me, which was emblazoned with the words “No Comment”. She didn’t have any business cards with her to give me her contact information, but suggested that I look her up

in PeeringDB, which is a service used by peering coordinators to publish and find details of peering information for organizations they're considering peering with [F-NANOG57:463].³⁶ Her negative response to my request mirrored the highly restrictive peering policy of her organization. Several of my interviewees, and attendees with whom I had conversations at NANOG, joked that most Tier 1 autonomous systems have a very simple one-word peering policy: "No!". In this case, the response I received was in no small part dictated by the economic position of her organization within the inter-domain routing system, and the larger landscape of the Internet and the telecommunications industry. At the other extreme, there are also some Tier 1 autonomous systems which have been known to voluntarily provide services to help resolve inter-domain routing issues, in the interests of the common good, rather than focusing on immediate economic concerns.³⁷ Commenting on why his organization has provided voluntary services to help resolve disputes, an interviewee told me:

We rendered assistance in a way that when a neighboring fire department needs mutual aid, you see somebody from another fire department go and take over the response mechanism there. If you remove the business equation from a lot of the discussions about the whole thing, the networks actually naturally want to connect and provide the best - the engineers want to provide the best technical solution that's available. You can't entirely remove the economic component, because it actually costs money to either buy the equipment or operate the equipment or whatever else. [I46:5]

An interviewee who is a long-time NANOG attendee framed the tension between immediate economic interests and open collaboration at NANOG well:

The larger providers have become more subject to "business strategic thinking", and therefore, they're less able to come to places like this and participate in it as an open forum for the exchange of ideas because they're less and less allowed to exchange their ideas. I think that that is an unfortunate disservice to the community and the Internet, as a whole, in the interests of fairly short-sighted commercial thinking that actually, in my opinion, doesn't provide the commercial benefits that it's expected to. [I25:6]

I experienced a range of responses to my requests for interviews, which were not quite as negative as the "No Comment" case detailed above, but did deviate from the norm of attendees who were willing to participate in a recorded interview. Some interviewees were hesitant about being recorded, and only agreed once I assured them that I would stop recording, or delete the recording immediately, if they became uncomfortable during the

³⁶See <https://www.peeringdb.com>, last retrieved Mar 20, 2014. PeeringDB is publicly accessible, and includes details such as an organization's peering coordinator contact information, ASN, approximate number of IP address prefixes, traffic levels and ratios, lists of IXPs at which they have a presence, and more.

³⁷While I do have concrete examples of such behavior, I have been asked by interviewees not to mention any details.

interview.³⁸ Another interviewee, who was a peering coordinator, refused to be recorded, but was willing to let me taking notes during our conversation, being of the opinion that nothing should be recorded at NANOG meetings [I45:1].

NANOG meetings are open to participation, but the ability to participate is patterned by a variety of factors. First, individual reputation permits a certain range of behaviors, and places - or relaxes - limits upon an individual's ability to transgress boundaries. Reputation is not simply a cumulative measure of lower or higher reputation, but must be understood in terms of the respect that an individual commands, and a collective acceptance or rejection of an individual's proclivities. Second, individual and collective perceptions of acceptable behavior construct boundaries which are not simply limits on behavior, but are also productive, in the sense that they privilege certain kinds of participation over others. For instance, the formalized limits upon commercial activity at NANOG enable a focus on technical content at NANOG conference sessions, in service of the values articulated in the NANOG mission statement. Third, organizational affiliation shapes the interests which an individual may express, occupying a spectrum between immediate economic interests, and coordination and collaboration activity in the service of the common good. These factors together determine the nature of participation, whether in conference sessions of various sorts, or in interpersonal interactions, shaping the degree and form of individuals' engagement in NANOG meetings.

6.4 Becoming a Network Administrator

The profession of network administration is an odd one. It requires highly skilled, knowledge-intensive labor, yet network administrators at even Tier 1 autonomous systems may have little formal training that qualifies them for their positions. The most widely accepted credentialing mechanisms are operated by networking equipment vendors, rather than by a third party, such as an academic institution; but these vendor-specific credentials are often considered insufficient in themselves as qualifications. Network administrators involved in managing critical components of the infrastructure of the Internet do so with skills and knowledge that are difficult to formally evaluate. How is it that the critical infrastructure of the Internet is operated by personnel with little by way of standardized formal credentials? How does someone become a network administrator? How is network administration recognizable as a profession? These are the questions I seek to address in this section.

6.4.1 Origin Stories

I began all of my interviews by asking interviewees for the story of how they came to be network administrators. With a few notable exceptions, all of those I spoke to fell into the profession without planning to; no-one grows up wanting to become a network administrator. Similarly, only a small proportion of those I spoke to had formal training in computer science.

³⁸I was never asked me to stop recording, or delete a recording, but there were several occasions during which interviewees explicitly told me that specific comments were to be off the record.

Whenever I brought this pattern up in interviews and conversations at NANOG meetings, those I spoke with confirmed my observation, often adding stories of others they knew who had, almost by accident, fallen into doing network administration.

As I documented in chapter 3, those involved with the ARPANET and NSFNET were largely computer science researchers, developing these projects as testbeds for computer networking technologies. As the network topology of these projects extended, they called for personnel to manage the day-to-day operation of network infrastructure, who themselves might not have a background in computer science. This was especially true at the peripheries of the NSFNET; there was a greater demand for computer science qualifications amongst personnel operating the NSFNET backbone. An interviewee with a computer science degree told me how she was hired out of graduate school by Merit to help build and operate the NSFNET backbone:

I got my Masters in '87, and then I decided I wanted to work for a while, because I've been having schooling for so long. So instead of going into a Ph.D. program I started to look for work, and then at the University of Michigan, at that time, I don't know if you know, there's an organization called Merit, which is a consortium of the state of Michigan, which takes care of networking of the state of Michigan. So at that time NSF put out the solicitation for the T1 NSFNET backbone. Merit, joined by IBM and MCI, got together, submitted a joint proposal to provide such a backbone. There were a lot of competing proposals, but Merit, MCI and IBM won the award. So they were hiring somebody, and I actually was the first one they hired to work on the NSFNET project. That's how I got into networking. [I5:2]

In comparison, consider this account from an interviewee who began his involvement with network operations at a university campus network in the same period:

I worked with biologists, entomologists actually, who were doing studies that were not amenable to the standard statistical approaches, so I started writing my own; that's how I got into computing. . . I moved into the networking group when the previous Apple Macintosh networking guy left the university, for Sweden actually. So I bought the book, and read it overnight, interviewed the next morning and got the job. And since then I've just stayed in the network group because I find it very enjoyable and quite challenging in ways that I like.

. . . During that time I participated in a couple of major upgrades to the campus backbone, where we replaced all the equipment and added whole new tiers of equipment, and I have been largely responsible for the purchase and configuration of all the campus border routers in that time, and that's really where my main focus is right now. My focus is largely with our external connectivity and our relationships outside campus. [I1:2]

This interviewee's career path was similar to that of many others I spoke to. Unlike the earlier network administrator hired to help operate the NSFNET backbone, he had no formal training that qualified him for his job. He picked skills up on the job, learning technologies as needed, until he came to occupy a senior position within the group responsible for managing the campus network, managing equipment and business arrangements for interconnections with autonomous systems outside the campus. As the Internet grew in the mid-to-late 1990s, the increasing demand for network operations personnel was satisfied in large part by people with qualifications in technical areas other than computer science. An interviewee who now works for a large hosting company told me how he became a network administrator:

Circumstance, happenstance, like everybody else. I was actually a graduate in health information science which is more systems analysis for hospitals. Graduated university, was looking for a profession backing it, but I wasn't really too happy with it. I was trying odd jobs and one of them was working in a call support of a regional ISP. Just moved up the ranks and ended up getting into systems analysis. Actually being a systems engineer for a while. Then the company [redacted], years ago, they needed a network guy. Nobody really applied, I was the only one. Ten years ago it wasn't that in vogue, I guess. [I21:1]

This interviewee came to network administration by "circumstance, happenstance, like everybody else." His education in a technical field may have helped his transition to network administration, but it was by no means a prerequisite, nor was it for many of his peers. More important, perhaps, is an inclination for working with technology, with a formal technical education being indicative of such an inclination, rather than determining it. Others I spoke with came into the profession of network administration without any formal education. An interviewee who is employed at a research lab for Tier 1 autonomous system began his career in the early years of the Internet:

So, I think I was like 16 or whatever, that was in '96, I was doing a lot of dialing into bulletin boards, and I got into Unix administration. I kind of got bored of the Unix side of things, I worked on when I was like 16, 17, 18, for a few ISPs. Just being in that environment, I learned how a service provider, small scale, kind of operates, running things from dial-up modem banks, all those authentication aspects - RADIUS - and Unix servers, that's how I got into it. So I kind of got bored of being on the systems side, so I moved into networks, worked for some small companies, mostly in Chicago, working in network operations centers, and moving on more in a design/engineering role. I've been doing design/architecture kind of things since early 2000. [I18:2]

Another interviewee began his involvement with network administration in his teenage years as well, albeit in the 1980s. His story was quite different, as he had to learn how to build and administer computer networks for very practical reasons:

From 1982 on, I had my first job working for someone not in my own family and I was working for the University of California press doing catalog layout, so they, as a publishing company, had lots of books. They would send a catalog out to book stores, quarterly usually, and at that time, layout of something like that meant having a typesetter, phototypeset columns of text and using a stat camera on photographs of the cover of the book and then using a hot wax roller to put wax on the back of everything and putting it up on a big sheet of board and sort of laying everything out very carefully.

Then in 1983, they got some Lisas in, it's a precursor to the MacIntosh and you could sort of see the direction things were going. Then in 1984, they got me a Macintosh to work on and it - I think they were thinking that it would increase productivity . . . Then I was doing desktop publishing. . . . There was a protocol, which I can't even remember the name of now, for on-the-fly replacement of an image in a layout document. So you'd have a placeholder in the layout document, you'd have the actual image on a file server, then you'd have the printer. You'd send the print job from the computer to the printer, and the printer would pull the actual images off the file server and do on-the-fly replacement. That all, of course, required a network and a file server and so I got into setting those up. Of course, it never worked right on the first try, which meant I had to learn PostScript programming . . . once you've got all that figured out, you're in much greater demand for running the PostScript, the \$100,000 printer than you are for running one of the \$2500 computers to do the layout. [I36:1-2]

This interviewee gradually shifted emphasis from publishing to computer networking, and went on to found one of the earliest ISPs in California. As these accounts indicate, neither age nor a lack of qualifications were a bar to entry into the network operations world in the 1980s and 1990s. This is not to say that computer scientists have had no role to play in network operations, but rather that the skills required for network administration are quite distinct from those taught in a traditional computer science degree.

As the Internet has grown, and customers have developed a substantially lower tolerance for problems with their Internet connections, entities operating autonomous systems have become more formal in their internal organization, and the prerequisites for entry into network administration work have increased. New entrants to the field are often, at a minimum, expected to have obtained vendor certifications. A young network administrator, who began work at a Tier 2 autonomous system in the 2011, told me how he qualified himself for his job:

When I was doing my [bachelor's degree in] electronics and communications engineering, towards my last year there were these two subjects that were more interesting, or really got my interest compared to the others. Those were computer communications and networking, and Internet communication engineering. . . . Subjects like this got me exposed to how packets are routed and what packets

are, how IPv4/IPv6 addresses and stuff. So it was something that was really new. All I knew about Internet is there is some kind of satellite communication going on. I never knew there was so much of undersea cables, and wireless, and wire-line, optical fiber communications and stuff. . . . That's when I decided I want to pursue a higher degree where they would teach me how routing/switching works and how data is transferred from one place to the other. That's how I got into University of Colorado at Boulder where they have a special Master's program for telecommunications. . . . The whole course is very, very focused towards getting Cisco certification so that you build a specific mark in industry. [I49:1-2]

Although this interviewee has a Master's degree in telecommunications, and vendor certifications from Cisco, he is surrounded by senior staff in his organization who lack his formal qualifications. He went on to tell me:

Even senior network engineers around me in my team don't have any kind of certifications. They have just done a community school kind of thing. Back in the 1990s they got into the NOC with [company name redacted]. They have 15, 20 years of experience. So they've seen how Internet has evolved. . . . they have learned on the job. [I49:6]

These are just a few of the many stories that I was told as I asked people how they came to be network administrators, but they are indicative of patterns which I observed across all of my interviews and conversations. Formal training and certification are more important to enter the field of network administration today, in comparison to the mid-to-late 1990s. Regardless of formal qualifications, or lack thereof, all of those I spoke with presented a frame of mind adapted towards the practical concerns of building and operating computer networks, rather than developing basic computer science research in computer networking. This is not to say that these two positions are mutually exclusive; senior network administrators will often present work at NANOG to form consensus around new approaches to operating networks, and use their experience to influence standards development at the IETF. Indeed, many organizations operating critical elements of Internet infrastructure (such as Tier 1 autonomous systems) find it to their advantage to have employees who are involved in network operations research. Although network administrators may sometimes be concerned with research into network operations, their primary concern is always with the practicalities of operating networks.

6.4.2 Learning to Route

As the previous section illustrates, a significant proportion of network administrators come to their profession without formal training in computer science. Lacking formal education in the technologies they had to work with, how did these individuals learn the craft of network administration? Many of my interviewees drew a distinction between network administrators involved in the everyday work of network operations, and computer scientists (whether

researchers at academic institutions, or employees of networking equipment vendors). As one of my interviewees put it:

Those guys often don't have much operational experience, they come from the research side, the protocol development side. They're very sharp, they know exactly what they're talking about, but it's not operations. [I1:14-15]

In drawing a line between “operational experience” and the “research side, the protocol development side”, network administration becomes more clearly defined as a profession which draws from experiential knowledge that falls outside the ambit of computer science, even though the technologies it is concerned with are drawn from computer science. This division was constructed gradually. In the ARPANET and the NSFNET period, the division between computer scientists and network administrators was less well-defined, as those developing computer networking technologies overlapped strongly with those managing the networks deployed using these technologies. During the early years of the Internet (in the mid-to-late 1990s), there was a high demand for network operations personnel, many of whom learned their trade on the job. A network administrator at a Tier 1 autonomous system told me how he learned to work with BGP:

I learned BGP by taking over a network for somebody else and sitting down and using the online Cisco help pretty extensively. That's how I got started and then I then used Cisco support pretty heavily to figure out the rest. . . . Basically Cisco was very much a participant in the education process. When I was getting frustrated with things and I couldn't make the routing work the way I thought it should work, I would send these guys [Cisco personnel] an email and say, “Hey, here's what I'm trying to do”, or “Here's what isn't happening that I need to make happen.” They would either teach me, they'd say, “Here, use these commands and it'll do this”, or they'd ask me some questions and a couple emails later we'd figure it out. Or the other half of the time I'd get an email back saying, “That's a known bug. Here's some software that'll fix that. Go to this location, download this, load it on your router, and that'll take care of that.” [I40:4]

This interviewee's experience mirrored that of many others who I spoke with. He learned the basics of BGP himself, with some support from personnel at Cisco, which was then (and continues to be) a dominant networking equipment vendor. Talking about the practical skills involved actual network operations, my interviewee continued, commenting on the role that NANOG played:

NANOG's not the only place to find those people [network administrators], but it was one of the main places where you met those people, where you got educated on how to do all the things that a national ISP should do. What should the filtering policy be for your customers? How do you make sure that the way you're running your network doesn't break the whole Internet, and things like that? NANOG was the main - was virtually the only source of that kind of

information when it comes right down to it. There was nowhere else to get that information back then. Yeah, there were only - at that time back in '95 there couldn't have been more than 100 guys that knew BGP well enough to even - in the whole world - to setup the interconnections between the networks. It was not intentionally an exclusive world, but it was very restricted just based on availability of skill and talent and experience. [I40:4]

NANOG was - and continues to be - an important space in which the practical experience of seasoned network administrators could be handed down to newcomers. This is as much in service of problems which require coordination and collaboration, as it is about gaining knowledge and understandings based on others' experience. An important set of issues that almost any NANOG attendee has to deal with, after all, relate to making sure that "your network doesn't break the whole Internet." Coordination and collaboration across organizations is eased by having a larger population of network administrators with common understandings of how they should build and operate their networks.

The interest that people have in passing on their experiences with practical issues in network administration at NANOG may be understood in terms of three factors. First, for the reputation that a presenter gains, as I discussed in earlier sections. Second, as a process of sense-making, working out particular issues in a public forum to come to common understandings of how to deal with them. Finally, as a matter of ensuring that network administrators in other organizations have sufficient skills and knowledge so that they don't make the kind of mistakes which may "break the whole Internet."

Many of my other interviewees had similar experiences in learning how to manage their networks. An interviewee told me how he had to learn on the job:

It was really difficult. I mean, I had no training. I learned you had to learn on the job. At that time it was a little bit easier just because our customer base wasn't that big, our traffic wasn't that big. If you made some mistakes, oh well, you can accomplish it. I started coming to NANOGs in 2002, met a whole bunch of folks and that helped. I learned a lot through just books, but then I learned a lot more talking to folks here. I haven't seen him around but there was a fellow by the name of [redacted] who helped out a lot. [redacted] helped me out a lot too. It's a very collaborative effort around here. [I21:3]

As this interviewee puts it, he was able to learn by making mistakes, since the impact of his mistakes was limited by the relatively small customer base of his organization at the time. Like others I spoke with, he gained basic knowledge from books, and then established relationships at NANOG meetings which helped him develop his practical skills. His experience of NANOG is of a "very collaborative effort", although it is worth remembering how early NANOG meetings were much edgier. This edginess has not been wholly tempered in the modern incarnation of NANOG; the ability to participate, and collaborate, is conditioned by a variety of factors, which I have already discussed. This extends to the NANOG mail list as well, which remains an important resource for asking for advice, and passing on experience.

An interviewee told me how important the NANOG mail list was to his learning process, and offered his perspective on the nature of conversation on the mail list:

A lot of it was online. A lot of it was observing NANOG at the time. The quality of conversation for NANOG has, to some extent, gone down. And like a lot of Internet-based lists, a lot of people are very intolerant of questions from new people. So you ask ignorant questions, you get flamed by many people. And if you are lucky, you will hear one useful answer from somebody and that will actually make all the difference. [I38:2]

NANOG meetings and the NANOG mail list offer a communal setting through which practical experience may be shared, and common understandings worked out. Practical experience is itself produced through the everyday activity of administering networks, and the coordination required for the work of establishing and maintaining interconnections between autonomous systems. Accordingly, it is productive to examine the ways in which skills are produced and transferred through the topological arrangements of autonomous systems.

A network administrator from a university campus network related how he learned BGP by setting up a border router for the campus in response to IP address spoofing attacks, learning BGP through the process of collaborating with personnel at the NSFNET regional network providing connectivity to his campus:

I used that [the IP address spoofing attack] as a reason to get emergency funding for the first campus' real border router, and then worked with a fellow at BARRNET,³⁹ [name redacted], who's been around for many years now ... and we set up the original BGP connection for the campus, as a real border router, and the whole point was to be able to prevent IP address spoofing. ... So we had to get routers that would actually allow us to filter based on source and destination IP addresses, our routers prior to that time did not. So that's how I started with BGP. [I1:3]

The NSFNET backbone network administrator I spoke with told me how she was involved with teaching personnel at the NSFNET regional networks:

... usually the upstream people, they've been doing BGP, they know exactly what needs to be done, so when they start to peer with somebody who has no exposure or experience with BGP, then they can just say, here's the sample configuration, take a look, then you can put it on your router and it should work. And how do these upstream people learn it? ... I worked with [name redacted], when he was at BARRNET, I was at NSFNET, I actually worked with him to have his BGP session up with us. I taught BGP, because I helped develop BGP, and I tested the first implementation of BGP. So, you know, you have the people who

³⁹The Bay Area Regional Research Network, the NSFNET regional network serving academic and research institutions in the San Francisco Bay Area.

originally know BGP really well, who've done a lot of work with BGP, and then they kind of just trickled down, you peer with people, and they're not really familiar with BGP, you kind of help them to get it up, you offer whatever means of help. [I5:13]

As these conversations make clear, the knowledge and skills required to maintain BGP sessions followed the topology of the NSFNET, flowing outwards from personnel at the core backbone network to personnel at the regional networks, and from them onwards to personnel at campus networks. The relationships formed in this process are multivalent, providing a transfer of skills and experience, but also serving as channels for coordination to maintain interconnections between autonomous systems, and helping to form a larger technical community, such as that instantiated through NANOG.

Personnel at peripheral autonomous systems may not even see themselves as part of the community which forms at NANOG, or the earlier NSFNET Regional-Techs meetings. The campus network administrator I interviewed went on to tell me how he didn't see campus network personnel as part of the Internet community at the time, although his perspective has since changed, as the campus diversified its upstream Internet connectivity:

So my perspective then would have been much more limited than it is now. Because we connected to BARRNET at the time, BARRNET took on the responsibility of connecting us to the outside world. It really was BARRNET which participated in the Internet community, for inter-domain routing, rather than the [redacted] campus, and this was true of all the campuses which connected through the regionals. Over time, the campus has got more and more involved, and like I said, BARRNET is no longer in existence. [I1:3-4]

This pattern of the production and transmission of experience along links in the topology of the inter-domain routing system continues to hold for the modern Internet. A significant difference between the modern Internet and the NSFNET period is that these topological links are now typically the consequence of economic arrangements. Accordingly, network administrators at upstream autonomous systems may sometimes be bound by support level agreements which determine the degree of engagement they may have with customers, although they will often still provide some technical advice, regardless of contractual limits. A network administrator managing BGP sessions with customers at a Tier 1 autonomous system described the different levels of engagement he had with customers:

For a vanilla BGP configuration, that's not a big deal, so we would be - we would help out as much as we could there. For the ones where we managed their router, we could log in and do it ourselves or whatever. The ones where it was the customer's router and stuff like that, we never actually got in and logged on the router or did anything. It was, "Here. I'll send you a config you can apply that I think will work." [I16:4]

For customers with simple needs - such as those which have only one upstream network provider - it is sufficient to send the customer a simple BGP configuration to get the BGP session set up. In other cases, this interviewee's organization contracted to manage the customer's equipment, so the customer had no need of personnel with knowledge of inter-domain routing. Finally, in more complex situations, this interviewee would send configurations which he thought might solve a customer's problem, but would not get involved further. In these situations - which are typically for multi-homed autonomous systems - it falls on the network administrator at the customer's autonomous system to tweak the configuration to fit their needs, in the process gaining experience with managing BGP configurations, guided by the configuration and advice provided by personnel at their upstream network providers. It is generally assumed that any multi-homed autonomous system must have personnel with sufficient knowledge to manage connectivity across multiple upstream autonomous systems, following their more complex topological position.

There are, obviously, differential levels of skill across autonomous systems at different positions in the topology of the inter-domain routing system. My interviewee offered his perspective on the development of skills amongst personnel at small-to-medium ISPs:

It's, you know, some of the folks that are still kinda learning how this all works, that level of operational experience isn't there yet. They don't have 10-15 years of experience operating an Internet network to say, "We did this before; we screwed up. We know what we need to not do here." I think that's probably where a lot of the value from things like NANOG actually comes is you get sort of newer providers that can go in and go to one of those things and learn from the more experienced folks that have been here. Yeah, there is that instability along the edges, but usually it's fairly contained in a small space. [I16:11-12]

The instabilities that this interviewee's organization has to deal with are "along the edges", at the peripheries of the inter-domain routing system. "Newer providers" occupying these peripheral positions can go to NANOG to gain the skills and knowledge required to reduce these instabilities, and to establish the relationships they need to move towards more central positions. The relationships formed between technical personnel for the purpose of establishing and maintaining interconnections between their autonomous systems are qualitatively different from the relationships which are formed and reinforced at NANOG. In the former case, relationships are relatively thin, defined almost entirely in terms of the technical details required for interconnection. In the latter case, as I have already discussed, relationships are substantially thicker, capturing details of a broader social world, and building in notions of trust and mutual reliance.

As an autonomous system diversifies its upstream and peer connectivity (complicating its topological position), it increases the uncertainty it faces in managing its inter-domain routing infrastructure, as it has to manage routing across multiple upstream and peer autonomous systems. As it takes on more customers, or provides more mission critical applications, it increases the risk it must deal with in case of failures in its inter-domain routing infrastructure. Increasing risks and uncertainties drive autonomous systems to engage more deeply

in regional and international technical communities to enable coordination and collaboration, and sharing of knowledge and skills, across organizational boundaries. The process of learning how to operate inter-domain routing infrastructure is by necessity a combination of individual acquisition of skills and experience, the establishment of inter-personal relationships for the coordination needed in everyday activity, and participation in broader technical communities - such as NANOG - to enable wider collaboration and to form common understandings of practice.

6.4.3 Institutionalizing Practice

Although the skills and knowledge required for network administration are worked out through practice, they have a long history of being formalized in documents and certifications. Here, I discuss the ways in which practice is codified, recognized and perpetuated through formal institutional arrangements.

As I showed in earlier chapters, the RFCs produced by the IETF are not limited to standards documents, but rather span a range of subject areas, including documentation of experience with particular standards, and BCP (Best Current Practice) documents which - as the name suggests - specify widely accepted notions of how particular networking technologies should be operated. Amongst the most commonly cited BCP documents at NANOG is BCP38, “Network Ingress Filtering” (Ferguson and Senie 2000), which specifies how autonomous systems should apply filters at their borders to ensure - to the degree possible - that BGP announcements originate from autonomous systems which are actually allocated the IP address blocks that they are announcing.

BCP38 is a “best current practice” in the sense that it is not a compatibility standard such as BGP which, if violated, will result in an inability to interconnect autonomous systems at a technical level. Rather, BCP38 codifies a practice that is widely accepted by network operations professionals to be necessary to support the stability of the inter-domain routing system by rejecting spurious BGP announcements. Once codified and published as a BCP, it becomes easy for network administrators to point to the BCP document as representing a global standard for the “right” way to manage a particular computer networking technology. For instance, during a presentation by a cybersecurity officer from the US Department of Homeland Security at NANOG56, an attendee suggested - to much applause - that the government should use its purchasing power to force BCP38 compliance by building these requirements into its contracts for transit and hosting [F-NANOG56-ARIN30:120].

Formal standards may also be established for the levels of skills and knowledge required to engage in the practice of network administration, in contrast to the informal mechanisms for the recognition of technical knowledge through reputation that I’ve already discussed. Major networking equipment vendors such as Cisco Systems and Juniper Networks offer vendor-specific certifications which have become industry standards, acting as basic entry requirements for many network administration jobs. A young network administrator who obtained Cisco certifications while in college presented his perspective on the role they play in hiring decisions:

When you are Cisco certified, you are going in for an interview, making the managers and interviewers know that you have this basic qualification done. So even if you are not experienced, you can still apply the concepts that you have learned, and you would be able to catch up quickly when you start working. [I49:6]

Cisco began offering certifications in 1993. The earliest certification it offered, the Cisco Certified Internetwork Expert (CCIE), continues to be offered today, and is amongst the highest level of certifications that Cisco offers. The CCIE is meant to test both theoretical and practical knowledge, featuring a written exam and separate lab exam, which requires the configuration and troubleshooting of actual equipment within a limited time. To remain valid, the CCIE must be recertified every 2 years.⁴⁰

The young network administrator I interviewed thoughtfully established a contrast between formal certifications, and the value of practical experience, as he spoke about the difference between his professional experience, and the academic work he did towards his degree and certifications:

This has been a question that keeps two sides of my brain fighting all the time, because what's happened is when I got out of school, I had my CCNA and CCNP.⁴¹ There was one more course that I missed because I took another course instead, which would get you a CCNP switch certification. . . . Then you become a CCNP without any experience. You are just a lab guy. So I stopped at the CCNP route because my seniors or my alumni said it's equally important to know the real world production environment rather than just knowing stuff.

Then I went into [redacted Tier 1 autonomous system] and saw there were so many people that didn't have the graduate degrees. They didn't have any certifications. But because of experience, they were there for five, ten years and they had great positions in the company and real knowledge. Coming from a lab environment to a production environment, I didn't know many things, because in a production environment, the design is different. There are different kinds of problems that come up. And you've never seen those problems in the lab, actually.

So I started thinking, let me see, my experience is more important than getting a certification. And so, I stopped studying and started analyzing more what a production network looks like and what it does. Suddenly, my juniors, when I was graduating, and even today, my juniors are CCIEs. They have no experience, and within just two years they got their CCNP certification. Then they get their

⁴⁰See <http://www.cisco.com/web/learning/certifications/expert/index.html>, last retrieved Mar 26, 2014.

⁴¹Cisco Certified Network Associate, and Cisco Certified Network Professional. The CCNA leads to the CCNP, which in turn leads to the CCIE. For details of Cisco's certification levels and areas, see <http://www.cisco.com/web/learning/certifications/index.html>, last retrieved Mar 26, 2014.

CCIE theory exam without even touching a router in the production network.
[I49:5]

The tension between vendor certifications and practical experience that this interviewee points to carries over to the larger NANOG community as well. A senior network administrator, who is a long-time NANOG attendee, and highly respected within the NANOG community offered his perspective on certifications, laughing:

Well, nowadays maybe it's [certification] a requirement to get through an HR screening. For a while I was handing out a... What did I call it? BIBCE - Building Internet Before Certifications Existed. I was handing out numbers to folks: "Here! You're 12!" [I50:14]

Experience and reputation always carry greater weight than certifications; yet, certifications are often necessary today to get the jobs that can lead to experience and reputation. In recent years, NANOG has taken on this tension by extending its training programs, and constructing new training formats, separate from the main conference, to provide training in operator-specific knowledge of network issues, rather than focusing on vendor-specific information on how to configure routers [F-NANOG56-ARIN30:52-53]. Presentations from NANOG board members involved in this project frame it quite clearly in terms of being "a way to pass on knowledge to the next generation of operators", and "a way to bring new people into NANOG".⁴²

These statements point to an interesting way in which NANOG understands itself as being constructed through a combination of knowledge and membership: to be a knowledgeable network administrator in the North American region, is to be a regular attendee at NANOG meetings. This is in part an assertion of the value that the social relationships formed at NANOG bring to the practice of network administration. It is also aimed at ensuring the continuity of NANOG. As I discussed earlier, concerns have been raised that the pool of regular NANOG attendees is aging, and that NANOG is having trouble attracting younger network administrators. The training offered by NANOG is intended to "bring new people into NANOG", and keep NANOG viable as a social space for the "next generation of operators".

In these respects, NANOG represents an institutionalization of the relationships and practices required to address the inter-organizational dependencies which are an integral component of the work of network administration, especially at larger autonomous systems. As individuals learn the craft of network administration through relationships formed in practice, those involved in the profession of network administration form a community of practice. NANOG provides a representation of a slice of this community who work at more topologically central autonomous systems.

⁴²From slide 17 of the presentation at the NANOG56 community meeting, available at <https://www.nanog.org/meetings/nanog56/presentations/Sunday/sun.member.opening.pdf>, last retrieved Mar 26, 2014.

As a profession, network administration has followed an interesting path towards institutionalization, melding the social relationships required for practice with technical knowledge acquired through individual study and experience, through training, and through relationships. As the profession of network administration institutionalizes further, it is likely that it will continue to maintain this balance between credentialing mechanisms, and practical experience and relationships, mediated through institutions such as NANOG.

6.5 Producing Interconnection

In the previous sections, I discussed the emergence of NANOG as a critical site for coordination and collaboration, and as a location through which the profession of network administration may be understood. In this section, I build on these understandings to examine the ways in which the relationships formed at NANOG contribute to the production and maintenance of the inter-domain routing system. My focus here is on how the value of trust in these relationships is leveraged to combat the problems of risk and uncertainty in the inter-domain routing system which I discussed at length in chapter 4. These trust relationships amongst network administrators are complicated by the economic interests of their employers. Accordingly, I will also examine the ways in which economic interests and trust relationships come into conflict, and are reconciled.

The distinction I draw between trust relationships and economic interests is an analytical one. In practice, these form a continuum of daily experience for network administrators. Many of those I spoke with articulated sophisticated economic theories which balanced - to different degrees - technical relationships to manage interdependence, with economic motivations for organizations operating autonomous systems. Even while they presented unified theories of this sort, they also drew a distinction between management and technical personnel. A senior network administrator put this quite succinctly:

We get along amazingly well at the technical level. The corporate sides, like the de-peering between Sprint and Cogent, or the Cogent-Level 3 de-peering before that, both of those were done by pointy-haired bosses, that are that type of people that were looking at the bottom line, and didn't really realize what havoc they were going to wreak for their own business. [I17:16]

The distinction between idealized technical interests, and idealized economic interests, plays out in a microcosm amongst regular attendees within NANOG. On the one hand, there are technical personnel who come to NANOG meetings to meet their counterparts at other organizations, and renew old relationships. On the other hand, there are peering coordinators who come to NANOG meetings expressly for the purpose of negotiating interconnections amongst autonomous systems with other peering coordinators. This is a simplistic breakdown: many peering coordinators have technical backgrounds, and many technical personnel have a hand in the business side of network operations. However, this is a useful distinction

to make for the purposes of this section. I will first discuss the ways in which trust relationships are leveraged in the day-to-day practice of network operations. I will then follow with a discussion of the business of peering, and the work of peering coordinators. Throughout, I will show how these different concerns - technical and economic - are linked, by the nature of interdependence in the inter-domain routing system.

6.5.1 Trust in Inter-Domain Routing

Recall that the technology which enables inter-domain routing, BGP, provides no mechanisms to evaluate the veracity of routing announcements. While autonomous systems may be able to more reliably evaluate routing announcements from their immediate customers, they can only trust that their peers are sending accurate information in routing announcements. In chapter 3, I argued that the “trusting” nature of BGP was the consequence of the strong trust relationships amongst the NSFNET technical community, and the hierarchical topological structure of the NSFNET which allowed for centralized control of routing via the NSFNET backbone. In the far more complex topology of the Internet’s inter-domain routing system, centralized control is no longer possible, making trust relationships critical to the stabilization of the inter-domain routing system.

I have already shown how thick relationships are formed amongst NANOG attendees. These are in part formed through a need for relationships to enable coordination and collaboration, and in part a historical outcome of the trusting environment of NSFNET technical communities from which NANOG was originally formed. Here, I will show how these relationships help stabilize the inter-domain routing system, and how the value of trust in these relationships is at the same time a consequence of the risks and uncertainties in the inter-domain routing system, and inter-organizational dependencies in network operations work more generally.

Commenting on the importance of interpersonal relationships for network operations, a senior network administrator related the coordination activity involved in the mitigation of the hijack of YouTube’s IP address space by the Pakistan Internet Exchange:⁴³

They’re pretty freaking essential. Like the YouTube hijacking was mitigated over IRC.⁴⁴ People don’t talk about it, but that’s how it was mitigated. The appropriate people were blamed over IRC. People had people’s home phone numbers who called them up to say like, “You need to stop accepting these routes” . . . and there’s all kinds of other things like that that happened. . . . There’s two kinds of things you need to do. One is try to fix something. The second is kinda engage in mitigation. In the case of like a hijacking or in the case of a bad peering or something, mitigation would evolve announcing lots of more specifics.⁴⁵

⁴³See chapter 4 for details of this incident.

⁴⁴Internet Relay Chat, a system providing text-based chat services.

⁴⁵A “more specific” is a more specific IP address prefix. Recall that BGP routers prefer more specific IP address prefixes when making routing decisions. A legitimate IP address prefix assignee can mitigate the

Well, what if your upstreams don't take the more specifics? Well, now you need to get them to emergency push prefix lists to you as a customer, and then we need to turn up additional adjacencies⁴⁶ in different places so that other people can populate these more specifics. In the idealized world of humans not being social animals, that's not really necessary, right? Because you'd have really, really high quality, low-friction, functioning contacts everywhere you needed to, but the reality - and this is maybe something interesting for the work you're doing - is that the problem with the Internet is people with whom you have no direct adjacency or any direct contractual or financial relationship can ruin your network. [I23:8]

This interviewee's response highlights a number of features of coordination, and the relationships required for coordination. First, that network administrators are always on call. This may be an active choice, like being available on an IRC channel. This may also be based on information shared in the formation of interpersonal relationships, such as home phone numbers. Second, that the actual work to be coordinated is in itself not a novel activity; there are widely understood responses to issues such as prefix hijacking ("announcing lots of more specifics") which provide a basis over which coordination may take place. Third, that there is a strong awareness within the world of network administration that interpersonal relationships are essential for everyday work, especially so in relation to infrastructural breakdowns. A network administrator needs good relationships with upstream autonomous systems, and other autonomous systems with which they may not have an economic relationship, in order to rapidly effect remedies to problems. Fourth, the awareness that these relationships are essential as a response to the fact that "people with whom you have no direct [relationship] . . . can ruin your network."

All of the interviewees I asked about these issues echoed, and expanded upon, the themes elaborated above. An interviewee, who was involved with setting up a research project to study the inter-domain routing system, told me of the challenges he faced in the mid-1990s in persuading network administrators at large autonomous systems to interconnect with his autonomous system:

Yeah, well in the early days it was hard, because I was working at a university, and I wasn't one of them, you know, one of those guys, so there was a trust thing that had to be set up. . . You know, in the early days it was all based on trust between me and the people who gave me views.⁴⁷ Personal trust, right. I'm not going to leak your routes, or advertise something stupid to your network, that breaks you. It was all trust. [I10:8]

This interviewee relates the problems he had breaking into the system of interpersonal trust relationships which characterizes the network administrators responsible for operat-

damage from a hijacking of their prefix by announcing more specific subprefixes of their assigned prefix.

⁴⁶An adjacency is a point of interconnection between two autonomous systems.

⁴⁷"Views" are snapshots of a router's RIB.

ing large autonomous systems. He also illustrates how trust is predicated upon technical competence: network administrators trust their counterparts in other autonomous systems not to make mistakes which may cause problems for their peers. A network administrator who was involved with network operations in the mid-to-late 1990s told me how she handled situations in which technical personnel at an autonomous system were perceived to have insufficient technical competence:

When the upstreams know these people are shaky, they put a filter on them, so then they cannot do too much damage. And then if you have people you trusted, you may not do filtering. At that time we already advocated that everybody do filters on their downstream customers. Even if the downstream is an ISP, people may say, these people have no BGP clue, so with no clue people, then you've got to filter it. [I5:14]

Unstable routing quickly becomes visible in the inter-domain routing system, since BGP works by relaying routing announcements from one autonomous system to another. When an autonomous system is seen to be engaging in the kinds of behavior which can cause stability issues for inter-domain routing at their upstream network providers, the easiest choice available is to place filters on their BGP sessions.⁴⁸ BGP announcements originating from autonomous systems employing personnel with “no BGP clue” must be filtered to minimize the effects of their lack of competence. This results in a topological arrangement in which large autonomous systems employ personnel who are competent, and trusted, since they relay BGP announcements to one another from other, more peripheral, autonomous systems. These large autonomous systems cannot filter BGP announcements from their peers when these peers are themselves only acting as intermediaries. In the absence of a centralized policing authority, the only choice left is to establish trust relationships across organizational boundaries for large autonomous systems, and hire personnel who are already trusted by their counterparts in other autonomous systems. A network administrator at a Tier 1 autonomous system told me how these relationships are used to enable ongoing conversations amongst personnel in different organizations:

There's lots of communication through backchannels, lots of unofficial communication. There's lots of things that nobody can officially talk about, but if we can all share information about it, we can make the Internet a better place. So there's lots of communities, and the only way that you gain admission into a community is to be a trusted individual. [I14:10]

This interviewee relates how the sharing of information across unofficial channels - amongst trusted individuals - can help “make the Internet a better place.” In this framing, the idea of the Internet as a project to be developed in concert, of acting “for the good of the Internet”, is an integral component of the system of inter-organizational relationships and unofficial

⁴⁸Filtering of BGP announcements originating from immediate peers and customers is generally considered to be a best practice, as documented in BCP38 (Ferguson and Senie 2000), which I discussed earlier.

communication flows. Network administrators come to be trusted through as their competence becomes visible to other trusted network administrators. When I asked this interviewee how he maintained his stature at NANOG, he was puzzled at first, and then responded:

So for me, that's not so much a problem. We're a Tier 1 ISP, and the people that we peer with are Tier 1 ISPs. You have to get to a sufficiently high clue factor to be a Tier 1 ISP. I can see how that could be a significant challenge to some of the Tier 2 ISPs. Yeah, I think that's not really a problem that we face, though. [I14:10]

Network administrators come to be trusted by means of their organizational affiliation, as much as by their individual reputation. There is a general assumption that a Tier 1 autonomous system will only hire technical personnel who have "sufficiently high clue factor". The trust relationships formed amongst technical personnel who perceive one another to be competent are essential to managing issues which originate from autonomous systems with which there is no direct relationship. An interviewee expanded on these concerns:

Yeah, in the ISP [world] it's pretty big because if you've got a good contact in your upstream for example, and you can at least talk to them, and get something done, or if there's nastiness emanating from one particular network, then it's good if you can talk to somebody because really, if you don't have good contacts, the chances of getting - influencing anything is pretty slim really, right. You can send an email off. You can ring them up, but if you're not a direct customer of theirs and not a direct supplier of theirs, they'll just tell you to bugger off basically, so yeah. It's kind of handy for that, if you know people, or if you know somebody who knows somebody. [I27:4]

As this interviewee points out, this kind of coordination calls for a set of trusted relationships, which must be navigated to find "somebody who knows somebody" who can help resolve a problem, or ease everyday administrative activity which calls for coordinated actions across different autonomous systems. Another interviewee related how it was sometimes not interpersonal relationships, but rather personnel responding to calls for help via the NANOG mail list, which helped address issues:

That was certainly one of the real educational things for the watching things like NANOG. With your upstreams, you at least have a phone number that you can call. You have some sort of support contract, right. You'll get some sort of feedback. Some of the upstreams are better than others. But almost universally, the majority of people that were more than one hop away from you didn't want to even talk to you. If you're not their customer, they don't necessarily want to help you. Depending on how helpful the immediate upstream is that is facing them, you may or may not get any feedback. But NANOG was very much sort of like the NOC of last resort. You could post a question there. You could usually find some friendly person that's actually working on the routing, and you don't

have to dig through the fact that you don't have any contracts with them and go through three layers of their support. [I38:4]

The established correct behavior for handling an issue is to contact the immediate upstream autonomous system which is the source of the problem, and rely on them to talk with their peer autonomous systems, and so on. However, this process is by no means reliable, since it can take time to relay concerns, and an appropriate response relies upon an appropriate relay of information, and handling of issues, at each stage in the process. When an issue cannot easily be resolved, this interviewee suggests, NANOG is the “NOC of last resort”, allowing petitioners to get directly in contact with personnel who can resolve their issues.

Complaints about having to cut through multiple layers of support in order to get things done are common at NANOG. At breakfast during the back-to-back NANOG56 and ARIN30 meetings, those I spoke with at my table wished that they could cut through organizational politics to get technical personnel talking to one another, so that they could easily get to someone with “the right combination of clue and access to routers” [NANOG56-ARIN30:826]. Economic concerns tend to keep flows of information and control across organizations closed; technical concerns tend to force these flows open. Network administrators must manage this tension across management and technical layers within their organizations, and within their own individual perspectives on the composition of economic and technical interests which make up the inter-domain routing system.

Social relationships which enable inter-organizational coordination are not the only means by which network administrators maintain the stability of their autonomous systems, and of the inter-domain routing system. These relationships are also leveraged for longer term processes of collaboration which result in tools and data to ease the everyday work of inter-organizational coordination in network operations. Many of these tools and data are hosted by a variety of non-profit organizations which support the work of network administrators around the world.⁴⁹ In the chapter 5, for instance, I introduced Team Cymru, which publishes lists of IP address blocks known to be used by email spam operations. These tools and data are to a degree a substitute for certain kinds of coordination work, and more generally function to augment all kinds of coordination for network operations. They cover a range of different functions, of which I will explore a few for purposes of illustration.

Some tools, such as the email spam IP address blocks published by Team Cymru, provide a centralized repository of data which would be difficult to construct purely through interpersonal relationships. Aggregating this data, and centralizing its management, provide benefits to personnel at all autonomous systems, making the Internet as a whole more stable.

A variety of projects are devoted to aggregating routing information published by BGP routers, with the aim of providing network administrators and researchers the facilities to understand inter-domain routing incidents as they occur, and to provide historical perspectives

⁴⁹There are also several for-profit organizations which offer similar kinds of services, as well as active monitoring of inter-domain routing information. I shall restrict my discussion here to non-profit organizations, in the interests of focusing on the collaborative aspects of network administration work.

on the growth of the inter-domain routing system. Route Views, hosted at the University of Oregon, is one of the best known of these projects, maintaining inter-domain routing data going back to 1997, collected from a variety of locations around the world.⁵⁰ The European RIR, RIPE, provides a similar service with its Routing Information Service project, which also collects routing information from locations around the world.⁵¹ Both of these projects collect data from IXPs at which many autonomous systems interconnect, providing situated perspectives on the inter-domain routing system. Network administrators may use this data to check how BGP announcements from their autonomous systems are propagated, as these projects provide a simple means to evaluate routing information from multiple global locations.

A similar kind of tool, called a Looking Glass, is offered by many autonomous systems. This allows network administrators to evaluate routing of their IP address prefixes within the autonomous system offering the Looking Glass service.⁵² These tools do not provide an exhaustive view on inter-domain routing. In fact, such an exhaustive view is not practicable to construct, since this would require a commitment from every autonomous system in the inter-domain routing system to publish routing information to a central repository. As I've already discussed, there are both economic and practical barriers to the construction of such a repository.

Other projects, such as the RIPE Atlas and the NLNOG Ring,⁵³ provide facilities for active network measurements. In both cases, network administrators volunteer to maintain probes within their autonomous systems which may then be used by other network administrators to run queries across the entire network of probes. These include services similar to Looking Glass, and running active measurements to estimate the actual routing and connectivity characteristics between autonomous systems. RIPE Atlas allows anyone to host a probe, and provides computing time across its network of probes based on how long an individual has been running a probe. The NLNOG Ring is much more restricted, only allowing participants who operate autonomous systems which maintain routers carrying default-free routing tables. As the website indicates, the NLNOG Ring is “based on mutual trust”; abuse of resources on the NLNOG Ring will result in an immediate ban for the individual involved in abuse. RIPE Atlas currently provides coverage of 1971 autonomous systems in 140 countries, while the NLNOG Ring is currently restricted to 248 autonomous systems in 45 countries. Both of these projects have been presented at NANOG in the interests of increasing participation in the North American region.

The CIDR Report provides an aggregate view on the state of the inter-domain routing system, including statistics on the number of prefixes and autonomous systems visible, autonomous systems which could reduce their announced prefixes through aggregation, possible

⁵⁰For more on the Route Views project, see <http://routeviews.org/>, last retrieved Mar 31, 2014.

⁵¹See <http://www.ripe.net/data-tools/stats/ris/>, last retrieved Mar 31, 2014.

⁵²For a list of publicly available Looking Glass servers, see www.lookingglass.org, last retrieved Mar 31, 2014.

⁵³For details, see <https://atlas.ripe.net> and <https://ring.nlnog.net>, last retrieved Mar 31, 2014. NLNOG is the Netherlands Network Operator Group, analogous to NANOG.

spurious route announcements, and more.⁵⁴ A summarized version of the report is automatically dispatched to the NANOG mail list, and various other network operations mail lists, on a regular basis. The intention in this case is to provide a picture of the overall health of the inter-domain routing system, and highlight autonomous systems responsible for certain adverse conditions. The individual currently responsible for maintaining the CIDR Report is employed by the RIR for the Asia-Pacific region, APNIC.

To ease direct coordination activity between network operations centers at different autonomous systems, the non-profit Packet Clearing House (PCH) operates the Inter-Network Operations Center Dial-By-ASN (INOC-DBA). The INOC-DBA system provides a closed Voice over IP (VoIP) telephone network which allows network administrators to dial one another directly by their autonomous system number.⁵⁵

These tools are, in a sense, a materialization of practices of network administration, intended to support the continuance of these practices. Broad intentions of openness are captured in tools such as Route Views, RIPE RIS and the CIDR Report. More narrowly scoped senses of openness - openness within a restricted community - are expressed in the NLNOG Ring, the RIPE Atlas and the INOC-DBA. These tools support different perspectives on the interpersonal relationships and community involved in the practice of network administration. From a functional perspective, these tools are anchors, and sometimes substitutes, for activity that is conducted through interpersonal relationships. These tools are also powerful symbols of how activity may be conducted “for the good of the Internet”, as individuals volunteer their time and expertise to maintain and operate these tools.

Access to these tools reflects the power that personnel gain from the position of their autonomous system within the topology of the inter-domain routing system. Some tools are publicly available, intended to provide data which anyone may use. Others are more restricted, available only to personnel with a certain status and affiliation, mirroring the manner in which participation is structured at NANOG, and the manner in which more central autonomous systems must “trust” one another’s BGP announcements, while placing filters on those from more peripheral “untrusted” autonomous systems. In short, a common set of structures - defined by an autonomous system’s topological position and a network administrator’s individual reputation - determine an individual’s ability to establish and leverage trust relationships for coordination and collaboration across organizational boundaries in the practice of network administration.

6.5.2 Negotiating Interconnections

Social relationships and practices are essential to maintaining the stability of the interconnections amongst autonomous systems from a technical perspective. As I will show, they are also of importance to understanding how peering coordinators make strategic decisions about interconnection with other autonomous systems. The value an autonomous system provides

⁵⁴See <http://www.cidr-report.org/as2.0/>, last retrieved Mar 24, 2014.

⁵⁵For details, see <https://www.pch.net/resources/papers/inoc-dba/docs/qanda.html>, last retrieved Mar 31, 2014.

to its customers, after all, is only in its ability to provide them access to the worldwide Internet. Such connectivity is possible only through interconnection, sometimes amongst autonomous systems competing for the very same customers. The business arrangements - “peering agreements” - involved in inter-domain routing are characterized by this combination of cooperation and competition. The resulting tension between openness and secrecy is an essential feature of the institutional and practical arrangements which support the inter-domain routing system, as I have demonstrated in this chapter and past chapters.

This tension was a well-recognized problem even in the early years of the Internet. Recall from chapter 3 that the NSF funded the creation of Network Access Points (NAPs) alongside NANOG and the RADb to help manage the transition away from the NSFNET to the Internet. The NAPs were meant to be neutral points at which autonomous systems could interconnect with one another. The organizations operating autonomous systems were to make their money from their customers, rather than from charging each other interconnection fees. In the early days of the Internet, settlement-free peering agreements were common, in which autonomous systems interconnect at no cost to one another.

This ideal was confounded by geography. Since there were so few NAPs at the time, autonomous systems sometimes had to lease long distance circuits to interconnect with one another at a remote location. A senior member of the NANOG community remembered the issues he faced in this regard when he founded an early ISP in California, as traffic from his ISP had to be carried to MAE-East⁵⁶ to interconnect with other ISPs in California:

There was this big split, right? Because 70 percent of the Internet was in California and about 10 percent was around Washington, D.C. About 20 percent was everything else. The problem was that the 10 percent in D.C. had these vastly higher profit margins because they were all selling to the government, whereas all the rest of were selling to scrappy little R&D labs in people’s basements and universities that didn’t have much budget and so forth. . . .so 70 percent of the Internet in California meant 49 percent of the traffic was going from California to D.C. and back to California, right, to go through the MAE. That was - I mean, on top of the fact that we had lower margins already, that was just a vast expense, right? That was at a time when most of the cross-country lines were just being upgraded from 56K to T1.⁵⁷ [I36:8-9]

MAE-East was built in 1992, through an agreement amongst commercial ISPs local to the Washington, D.C., region. Private ISPs provided limited IP network services well before the NSFNET was privatized. MAE-East was later funded by the NSF, along with three more NAPs. Of these, the Californian NAP, MAE-West, was only built in 1995. In the interim, a group of Californian network operators collaboratively created Packet Clearing

⁵⁶Metropolitan Area Ethernet East, a NAP in the Washington, D.C., area.

⁵⁷This is an indication of the data rates on these lines. “56K” refers to 56 Kbps, and “T1” indicates a line which can carry 1.544 Mbps.

House (PCH) in 1993 as a vehicle to construct local IXPs,⁵⁸ which offered them substantial savings over the cost of long distance circuits to interconnect their autonomous systems at MAE-East. My interviewee related the origins and organization of PCH:

A bunch of us who had ISPs in California got together and got PCH started and built five more exchanges [IXPs] up and down the West Coast, which kinda took care of that problem. PCH at that time was just a project; it was not incorporated. It didn't have its own budget or staff. Everything was seconded over from our for-profit companies. Then in, I think, '97, we incorporated as a non-profit. [I36:9]

PCH was a regional collaboration which provided economic advantages to all involved. However, it required substantial commitments of trust amongst the organizations contributing to it. As my interviewee noted, resources were provided by contributing organizations as needed. These trust relationships are quite apparent in the manner in which PCH was created, as an informal collaborative enterprise supported by voluntary contributions, which was not incorporated until 1997.

There are currently more than 350 IXPs in the world today, of which more than half are in the USA and Europe (Weller and Woodcock 2013:14). These are sometimes provided as for-profit services, and sometimes as operated as non-profit associations of ISPs, with variations across this spectrum of possibilities, which I will explore further in the next chapter. IXPs provide the immediate benefit of avoiding “tromboning”, keeping traffic local, rather than relying on a remote location to interconnect geographically proximate autonomous systems. Once an autonomous system has invested in infrastructure to connect to an IXP, it has an incentive to interconnect with as many other IXP participants as possible. As one of my interviewees put it:

There are hundreds of people that have built networks that show up to exchange points and the popular data centers centers in order to reduce their transit bill. . . . Once you've made the step of - you've made that investment, you need to peer with everyone you can at that location. You need to maximize the value that you are getting for that expense that you've outlaid to build that network. Because there are hundreds of people in the same boat as you, you can get peering on all kinds of people. [I40:14]

At the simplest level, an IXP provides technical facilities for interconnection amongst many autonomous systems, leaving the bilateral agreements for interconnection to the participating autonomous systems. Some IXPs which are managed as non-profit associations provide multilateral peering agreements, in which all participating autonomous systems agree to interconnect under common terms. Multilateral peering agreements offer an additional benefit to autonomous systems, in that they are able to interconnect with many peers under

⁵⁸IXPs and NAPs are functionally equivalent; NAP is just the term which was used by NSF at the time. I will prefer the modern term IXP in my narrative, except in reference to the historical form of the NAP.

a single agreement and set of technical arrangements, rather than having to negotiate and maintain multiple bilateral peering agreements.

The largest multilateral peering agreement, at the Frankfurt IXP, involves 271 participants (Weller and Woodcock 2013:65). The largest IXPs have several hundred participating autonomous systems.⁵⁹ IXPs are by no means the only mechanism for interconnection of autonomous systems. Under certain conditions, autonomous systems may interconnect with one another at private peering points; such arrangements are especially common amongst Tier 1 autonomous systems which establish many geographically distributed points of interconnection with one another. Once autonomous systems establish an interconnection at an IXP, they may later engage private peering at other locations as they each expand their business. In all of these cases, relationships amongst peering coordinators are critical to producing and maintaining peering arrangements.

Informal arrangements were integral to the business of peering in the early Internet, although these dynamics have since evolved. A senior network administrator told me how the process of forming peering agreements has changed over the years:

I guess in my case I sort of grew up in a community that had already been there for a while. The problem is certainly that in the last 10 years or so, if not a little bit longer, things have been transitioning to a much more commercial basis. A lot of the ways that people get things have done have changed. It used to be that 10 years ago, if people needed to set up peering, they would be wandering the halls of NANOG looking for people in appropriate networks. There would be handshake deals in the hall that, in some cases, proxy 4 million dollars worth of traffic. Such things are certainly less than needful these days. There are things like the peering database [PeeringDB] that help arrange those sorts of meetings. Bottom lines are much more strictly controlled for the networking stuff. The days of cowboy operators running . . . are a lot harder to do. The networks are much bigger, and, therefore, the people that are involved tend to be more business types. [I38:1-2]

In the early years of the Internet, in the mid-to-late 1990s, handshake agreements for interconnection were common, as this interviewee points out. As autonomous systems transitioned away from “the days of cowboy operators”, and Internet connectivity has become a critical service, the interconnection of autonomous systems has come to be governed to a greater degree by business concerns. This is not to say that markets for interconnection have gradually made moves to disembody themselves from social relations, but rather that the form of embeddedness of markets has changed to be increasingly conditioned by the topological position of autonomous systems in the inter-domain routing system. Social relations remain central to the construction of network interconnection markets, but have changed from a relatively flat organization to a system in which a hierarchy of positions becomes

⁵⁹For a directory of IXPs, see <https://prefix.pch.net/applications/ixpdir/>, last retrieved Apr 2, 2014.

clear, conditioned in part by an individual's reputation, but also by the topological position of the autonomous system they represent.

Larger autonomous systems gradually began introducing a greater range of conditions for settlement-free peering, from the number of geographically distributed locations for interconnection, to the amount of traffic exchanged, to the ratio of traffic exchange, and most recently to measurement in terms of bit-miles [I40:8-10].⁶⁰ NANOG was perhaps the most important gathering in the North American region at which network operators could find potential peers to interconnect with. This is still true to a degree, but in different ways than it was in the early years of the Internet. Another interviewee expanded on these changes:

NANOG's evolved a lot. It used to be a fairly small community of people, and all of the Tier 1 ISPs were represented there with a fairly good group of folks. Now, a lot of NANOG - I mean there's still interesting things going on technically, and from that perspective, it's still useful to be there. The interaction is more between different companies, Tier 2 ISPs, things like that, where they're arranging peering across the peering fabrics and the Equinixes of the world and things like that.⁶¹

Where that's kind of - it's a sticky thing for most Tier 1s because they have very strict settlement-free interconnect policies that govern who they peer with and who they don't. You get a lot of people who would like to have peering and don't qualify for one reason or the other. It's not trying to be elitist or anything; it's simply that's the nature of the business. That portion of it, there's not as much value for companies like [redacted Tier 1 autonomous system], but being there still offers an awful lot of networking opportunities, and there are still a lot of potential customers that will stop and say, "Hey, I'm thinking about doing this or that," so from that perspective it's useful. [I16:7-8]

Tier 1 autonomous systems form a closed club into which it can be difficult to gain entry, due to the stringent requirements that Tier 1 autonomous systems establish for settlement-free peering.⁶² As this interviewee notes, NANOG meetings are of greater utility to Tier 2 autonomous systems seeking peers, in the interests of reducing their transit costs through Tier 1 autonomous systems. Some of these meetings occur by happenstance in the hallways at NANOG meetings, as peering coordinators identify one another by the green dot on their badges. Meetings to negotiate peering are also formally scheduled, as peering coordinators evaluate one another's autonomous systems using the data available in PeeringDB. The data captured in PeeringDB is indicative of processes that have developed over time, as peering coordinators have progressively refined their common understanding of the factors involved

⁶⁰See chapter 4 for a discussion of these factors.

⁶¹A "peering fabric" is term which refers to the network infrastructure required for peering, typically in the context of an IXP. Equinix is a corporation which operates data centers and for-profit IXP services. See <http://www.equinix.com>, last retrieved Apr 1, 2014.

⁶²See chapter 4 for a discussion of the requirements that Tier 1 autonomous systems place on settlement-free peering agreements.

in negotiating peering agreements. Speaking about the the value of the interactions amongst peering coordinators at NANOG meetings, a peering coordinator for a large content provider told me:

Yeah, it makes it very useful. Peers tend to have defined processes and these processes are results of everybody coming to these meetings and then coming up. Hey, PeeringDB was a result of these meetings. Understanding peering, understanding resources that are involved in the peering infrastructure certainly helps fast track it. If you just turn on the peering and then try calling up their NOC, you're in the ticket hell. You're not going to get anywhere. Because everybody's aware, okay, there's usually a side step here to bring up peering. [I21:5-6]

This comment captures both the value of ongoing interactions in defining processes and tools for peering, and the value of interpersonal relationships established between peering coordinators to ease the establishment of peering sessions. Interpersonal relationships allow peering coordinators to bypass formal organizational structures and processes to set up peering sessions. In fact, as this interviewee suggests, "everybody's aware" that interpersonal relationships are expected to be leveraged to facilitate these processes. I asked a senior network administrator about the kinds of relationships that are required across corporate boundaries to maintain the inter-domain routing system, and he immediately responded by talking about the work of peering coordinators:

Yeah, the across corporate borders thing, where I see that being most visible in some ways is in the peering community. There are fewer well connected peering coordinators than there are transit ASs that peer, right? When a transit AS needs to grow quickly, they have to attract a peering coordinator who knows the other peering coordinators, and the set of those is not quite big enough to go around. It's kind of like reverse musical chairs or something. Good peering coordinators move from company to company to company doing the same job at each company, benefiting each company equally, but with no loyalty to a company, because their loyalty is to the group of peering coordinators who are the other people who make what they do possible, right?

There will always be another company that needs more peering, but the relationships that allow them to pick up the phone and say "I'm now working for Backbone X and we need to peer in these 15 locations, let's get that turned up" - that relationship is not replaceable. That relationship is with an individual who will be at another company themselves next week, right? [I36:15-16]

In this interviewee's perspective, the role of peering coordinator is not simply one into which any qualified person can be placed. Rather, the role is defined by the relationships that the person who fills the role has developed with other people in similar roles. The work of a peering coordinator is - as the title suggests - rooted in *coordination*, which requires the

careful cultivation of relationships with other peering coordinators. These relationships are tied to individuals, not organizations. Another interviewee echoed this perspective:

There's people who have experience in the industry, who understand the nuances of peering negotiation, interconnection between service providers. If you're a up and coming provider and you're trying to get kind of connection to a whole bunch of people, knowing the right people to talk to, and because of your reputation and your history, will ease the business. . . . There's companies that will say, "Yeah, you have these relationships, and we'd like to take advantage of them." They come to whatever agreement it is on salary to decide that. Salary and compensation and travel and whatever that is, and there's what some people call the fine-dining peering circuit where they go to all these exotic locations, and they have meetings every couple of months. [I46:10]

Relationships amongst peering coordinators are characterized by common understandings of how peering works and is to be negotiated, and common knowledge of the histories and relative positions of organizations operating autonomous systems. Anyone with a green dot on their badge at a NANOG meeting is expected to be able to talk about these subjects in relation to their employer, as with the case of my acquaintance from the story with which I opened this chapter. Peering coordinators gather in person at a variety of meetings - including NANOG meetings - to conduct peering negotiations and renew their relationships. These two activities are by no means mutually exclusive; a peering coordinator from a Tier 1 autonomous system told me that peering coordinators won't respond to email unless it's from someone they've already met in person, but that once they have formed an acquaintance, "you can hardly blow off a guy you've had a beer with!" [I45:17]. The relationships that peering coordinators form are not limited to their functional outcomes; they are thick relationships, like the other relationships formed at NANOG meetings that I've already discussed. A peering coordinator from a Tier 1 autonomous system spoke about the importance of forming relationships by meeting in person, to evaluate how trustworthy other peering coordinators might be as individuals:

It's important to come here to build that friendship, too. Quite frankly, at the beginning you don't know who you can be more open with, whom you have to be reserved. You have to look the person in the eyes. Through their track record, you kind of know, "OK. I can work this person this way. I can work with this person that way." That will also help you to be more effective.

For example, some people like to play games, and you have to play that game. . . . Some people do not and you can cut to the chase. And they appreciate it and things can be done much faster. . . . So you have to find the right attitude or strategy to work with different companies, different people. [I47-48:6]

Although relationships are held by individuals, organizational affiliation does shape the work of a peering coordinator. Different organizations have different perspectives on how

to manage the role of a peering coordinator. Some may align their peering coordinators with engineering, while others may align peering coordinators with sales. These choices have material outcomes; a peering coordinator from a Tier 1 autonomous system told me how she perceived an alignment with engineering to facilitate a “more neutral position, to have a more global view” [I47-48:4]. Those aligned with sales, in contrast, might be more focused on their organization’s economic interests, rather than technical interests which require greater considerations of interdependence.

As autonomous systems grow, the function of a peering coordinator may be implemented in institutional, rather than individual, form. A senior executive from a Tier 1 autonomous system told me how peering coordinators in his organization have substantially less power than those at other organizations, since the business of peering is part of organizational structure and process, rather than led by an individual:

Our peering coordinators at [redacted Tier 1 autonomous system] are actually almost more provisioners than they are business people. We’ve also had a steering committee where peering requests are brought to. The peering coordinators, kind of their power has been vastly reduced. Our peering coordinators here do not have the ability to, at least not very easily, do handshake deals with other peers to get things done. . . . If you have a sufficient amount of bureaucracy in place, I guess is the way to put it, you can reduce the power of the peering coordinator tremendously. Most networks don’t have that level of bureaucracy unless you’re talking about the largest telcos and backbones out there. Most of them don’t have that. [I40:13-14]

The role of a peering coordinator is shaped by organizational culture and structure, and also by the topological position of an autonomous system within the inter-domain routing system. For instance, a peering coordinator at a Tier 2 autonomous system may be more concerned with setting up new peering relationships, while a peering coordinator at a Tier 1 autonomous system may be more concerned with expanding existing relationships, and limiting new relationships to contracts for transit. There is a tension between trying to establish settlement-free peering, and signing up paying transit customers, which can influence the strategy that autonomous systems adopt in managing the capacity - the aggregate bandwidth - of their peering sessions. This tension has to be handled carefully; a senior network administrator told me how a reluctance to increase capacity on peering sessions can create bad blood:

I hear a lot of people talk about like we’re gonna have a quarterly capacity review. I hear a lot of people point out specific providers that they hate because they won’t do good capacity planning on their peering, because they want everybody to buy transit. [I23:12]

These kinds of strategies are particularly visible when a Tier 2 autonomous system is trying to grow towards Tier 1 status. I heard these concerns first hand during a conversation

I had at NANOG58 with two network administrators who were employees of different Tier 2 autonomous systems. They were not themselves peering coordinators, but were involved in setting up peering sessions, and had coordinated with one another for this purpose on several occasions. One of them spoke about how his organization was trying to grow towards being a Tier 1 autonomous system, and in doing so had to become more selective about peering. He talked about how he regretted having to turn away people who were ready to peer. The other network administrator nodded in agreement, agreeing that they needed to have fewer peers, and more transit customers, in order to grow towards Tier 1 status. He suggested that they change their peering policy from “open” to “selective” in PeeringDB to indicate their new direction [F-NANOG58:136]. As an autonomous system grows to Tier 2 status, it must engage in settlement-free peering with as many other autonomous systems as possible, to reduce its transit costs over Tier 1 autonomous systems. In order to grow to Tier 1 status, however, an autonomous system must reduce its settlement-free peerings, and instead focus on increasing its transit customers. This move will increase its customer traffic, moving it towards the aggregate bandwidth levels required by Tier 1 autonomous systems to be considered for settlement-free peering.

Peering coordinators are often intimately involved with this kind of strategic thinking. All autonomous systems - especially ISPs and large content providers - depend upon an appropriate mix of settlement-free peering and paid transit links to manage their costs, and plan their future growth and evolution. A senior network administrator gave me his perspective on the value of a good peering coordinator:

... being able to keep up with a good peering coordinator who's working hard is a lot of budget and a lot of other people doing network planning and engineering and implementation and deployment and racking and stacking and all that, right? A good peering coordinator - one good peering coordinator can probably keep 500 to 2000 other people busy, right, doing network build out. It's never seen that way, right? It's always seen as the senior network engineer decides where we gonna go, right, and then everybody follows their directions and the peering coordinator comes along behind and makes it do something, right? That's true, but if the peering coordinator isn't there or isn't able to do that job, eventually you run out of money for everything else, because the return on investment is too low. [I36:16]

A senior network engineer may lead the process of capacity planning for an autonomous system, making decisions about where to extend network infrastructure; but it is the peering coordinator who makes the agreements for interconnection which lend value to that infrastructure. Equally, it is sometimes the peering coordinator who helps form strategy about where network infrastructure needs to go. A peering coordinator is able to inform these decisions through knowledge of their autonomous system's current peering agreements, and how those fit within the broader landscape of peering in the inter-domain routing system. Peering coordinators arrive at this knowledge through constant engagement with other peering coordinators, often in face-to-face interactions at NANOG meetings and other venues,

since the details they share with one another are sometimes sensitive corporate information. A peering coordinator for a Tier 1 autonomous system told me how she perceived meetings such as NANOG to be important for sharing information:

A lot of businesses are under NDAs, so sometimes they don't feel comfortable talking to you over the phone or email. So you have to meet in person. Usually, in an environment like this they can be frank. They can tell you things that they cannot tell you via email or phone call. . . . Usually, when something gets done, it's something that both want. So you just have to bridge that gap to get something that both want. [I47-48:5-6]

Peering coordinators are involved in a variety of face-to-face interactions at NANOG meetings, some by appointment, some by chance. However, not all interactions between peering coordinators are one-on-one. NANOG meetings provide a special session at which peering coordinators gather, called the Peering BoF. Unlike other sessions at NANOG, the Peering BoF is never recorded or broadcast, a convention which attendees are reminded of before the session begins. Accordingly, I will not report on the specific content of these sessions, although I will describe their structure and format in general terms. The intention behind the prohibition on recording is to ensure that those in the room can have conversations freely - potentially sharing sensitive corporate information - in the interests of continuing to develop common understandings of peering, and forming connections that they might not otherwise have been able to.

Although there are limits on recording at the Peering BoF, there are no limits on participation. It is open to any NANOG attendee, not just peering coordinators. Unlike other sessions at NANOG meetings, chairs are typically arranged in concentric circles, with everyone facing towards the center of the room. There are microphones placed in aisles cutting through the circles for attendees to engage in comment and conversation, and a projector screen is provided to one side of the circle of chairs to display presentations.

Like the NANOG meeting as a whole, there is a division between frequent participants in the Peering BoF, and others who are finding their way in. Those sitting in the more central circles of chairs are constantly engaged in conversation with their immediate neighbors, and across the circle, and clearly know each other well. These interactions are quite boisterous, and remain so throughout the Peering BoF. Those who sit in the more peripheral circles - such as myself - are there to observe, and potentially find their way into relationships in the more central circles. At one of the Peering BoFs I attended, for instance, my neighbor - who was employed by an IXP provider - was attending just his second NANOG meeting.

In a manner of speaking, the arrangement and population of chairs is legitimate peripheral participation made concrete: there is no absolutely central position, only the possibility of moving towards a more central position through continued engagement and the formation of relationships within the community of practice of peering coordinators. This organization also mirrors, to a degree, the topology of the inter-domain routing system; those who are frequent attendees, sitting in the more central circles, are likely involved with Tier 1 au-

onomous systems, large IXPs, or large Tier 2 autonomous systems, while those who sit in the peripheral circles likely represent smaller Tier 2 autonomous systems, or emerging IXPs.

The Peering BoF consists of a limited number of formal presentations, where presenters discuss subjects ranging from technologies involved in the practice of peering, to the economics of peering, to the overall state of the market for peering in North America. There are frequent interjections during these presentations from the more central circles, although attendees will also wait to queue up at microphones to engage in comments and questions after a presentation is complete. Attendees commonly address one another during the presentations, just as much as they address the presenter. The interactions at the Peering BoF often make clear the opinions attendees hold of one another, in very public ways; albeit limited in visibility to those physically present.

Following the formal presentations, there are presentations from IXPs, who provide updates on their current status and future plans, with the intention of attracting more peering sessions from the peering coordinators in the room. The success of an IXP, after all, is measured by the number of autonomous systems which peer at it, and the amount of traffic they exchange; the value of peering at an IXP increases with the number of participating autonomous systems. Finally, there are “peering personals”, during which any attendee may stand up and provide a standardized set of information about their autonomous system, with the intention of attracting new peers.

The Peering BoF is part a networking event at which peering coordinators may discover potential new peers and sites for peering, part a sense-making event at which understandings of peering are worked out, and part a social gathering of friends. The Peering BoF at NANOG meetings is not the only such event at which peering coordinators meet. There are several other similar events, some of which proceed over multiple days, such as the Global Peering Forum (GPF).⁶³ Attendance at the GPF is restricted to peering coordinators. GPF takes place at locations ranging from resorts to cruise ships, with substantial sponsorship from IXP providers, who stand to gain from increased peering activity, as providers of interconnection infrastructure. This sponsorship is also due to the fact that GPF was created as an amalgamation of peering coordinators’ meetings that individual IXPs hosted. As the number of such meetings increased, it became difficult for peering coordinators to attend all of them, which resulted in a collaboration amongst IXPs to bring these meetings into a single venue [I45:15].

Peering coordinators’ meetings, such as the NANOG Peering BoF, and the GPF, operate in a limited closedness - prohibiting recording, or vetting attendees - in order to enable openness of discussion amongst attendees. Unlike the other social settings I examined in this chapter, and the last, openness cannot be understood simply in terms of the work that goes into producing openness. Rather, openness must be placed in context: peering coordinators need to be able to share sensitive corporate information in order to arrive at peering agreements, and make sense of the practice and market of peering as a whole. They can only do so under conditions in which they can assume that the information will not be

⁶³See <http://www.peeringforum.com>, last retrieved Apr 3, 2014.

shared outside the context within which it was revealed. Context in this case may be an interpersonal relationship, a peering coordinators' meeting, or the network of relationships amongst peering coordinators (and subsets thereof). This network of relationships is not strictly limited to peering coordinators, as it extends to encompass other trusted individuals who attend NANOG meetings, and other similar meetings. As people began to recognize me at NANOG, for instance, I was occasionally told confidential stories, always with the express understanding that they were not to be shared.

To become a peering coordinator is to become part of the network of relationships amongst peering coordinators; and in the process become part of this community of practice. As I've already noted, the practice of being a peering coordinator requires the maintenance of relationships with other peering coordinators, through which the practice is carried out, and made sense of.

It can be difficult to become part of this community. One of my interviewees, who has been involved with peering at several Tier 1 autonomous systems and IXPs, told me how it can take years to break into the community of peering coordinators. He was mentored by four different people who wanted to see his IXP succeed. Their approach was to introduce him to the community of peering coordinators, and then leave him to fend for himself, which he likened to being "thrown to a pack of wolves" [I45:5].

Others I spoke with had an easier time entering the community of peering coordinators. A peering coordinator at a Tier 1 autonomous system reflected on her process of assimilation into the community, and how she relied on two generations of her predecessors to smooth the way, along with a senior member of the community who saw value in mentoring her:

I was lucky that, when I started this position, my predecessor was still here. The predecessor before him was still there. So they make a lot of inroads for me. So the transition was pretty smooth. Of course he already knew a lot of people. We kind of partnered for a while before I came . . . We still, in a way, partner. But I came to gain more and more responsibilities. So the transition for me was very smooth.

And also, sometimes it's depending on who you bump into. I remember when I first became a peering coordinator. The very first NANOG that I went to, I bumped into this person and he happened to be someone that knows everyone in this community. So I just hung around with him in the hallway for an hour. So you kind of know a lot of people. And he's so kind to introduce you to all the people that he knows. Of course he also has his agenda. Because by helping you to succeed in this business, he has things that he needs from people doing peerings. This is a win/win situation. . . . I just happened to bump into him in the hallway. This was just luck. It's not something I can teach or something that you can prepare for. Sometimes luck has to play a part of it.

The bottom line is, in order to break the ice, time, time, time, time. You have to spend time in NANOG. One thing I do find is that the more you come to NANOG,

the more you will be able to be accepted by the others. This is something that I always tell the newcomers, is oftentimes you come here and you don't know anyone. People don't know you. But you come back again, they'll say, "Oh, I know this face before." And they will start trying to learn your name, try to remember you.

But if you disappear, they don't want to spend effort that someone may disappear next time. But you come here like a piece of furniture. Every time you come, they will see you like a table or a chair. You will be sitting there. They will say, "OK. Let me see whether this chair or table is comfortable." It's become worthwhile to know you. A lot of people here are very busy. They just don't have the time. Even if they try to do it, either they never bump into you or they just don't bother knowing new people until they say, "Oh, this looks like a familiar face." [I47-48:7-8]

As this interviewee points out, it is not sufficient simply to establish relationships. It is critical to attend NANOG and meetings of peering coordinators regularly, maintaining a presence "like a piece of furniture". Relationships amongst peering coordinators depend upon recognition and familiarity within the shared physical spaces of peering coordinators' meetings. As this discussion suggests, there is a high degree of trust required amongst peering coordinators. They share sensitive corporate information with one another, and often close peering agreements simply with a handshake deal. In a recent study of peering agreements, 99.5% of peering agreements surveyed were concluded without any written contract, thanks to commonly understood "rules of the game" (Weller and Woodcock 2013). A peering coordinator with a large Tier 2 autonomous system told me how these agreements worked in practice:

And settlement-free peering, you are not really exchanging any money, obviously. So contracts are not really...I won't say necessary. A lot of people, and my company is one of them, legal departments get kind of worried about, "Well, what happens if they are able to get into our network somehow?" Well, one of the things about this community is if you do something wrong, everybody knows about it. If you do something that's kind of an unwritten rule, like tunneling customers, it's an unwritten rule. It will be known. Somebody will tell somebody. Everybody will tell everybody else and nobody will do business with you after that. It's happened a lot. Not a lot, but the instance, if it's a bigger company, yeah, you'll hear about it. You'll read about it. It will get posted everywhere on the Internet. Let's face it. That's what we are, we're the Internet. So it gets known. [I42:3]

Those who violate the "rules of the game" face public censure, and possible ostracism, both as individuals, and as organizations. Cogent Communications has faced these kinds of

issues for several years now, as it engaged in peering disputes with a variety of ISPs.⁶⁴ However, Cogent's attendance at NANOG has dropped off in recent years, possibly as a result of the strategies Cogent has employed which have resulted in peering disputes. NANOG attendance records indicate that Cogent employees attended NANOG regularly from NANOG19 (in June 2000) to NANOG49 (in June 2010), after which they attended no meetings until NANOG60 (in February 2014). In 2010, Cogent was removed from PeeringDB, since Cogent sales people were using contact information in PeeringDB to spam personnel at other autonomous systems with sales offers. All pages on PeeringDB include a footer which clearly states, "NOTE: Sending Unsolicited Commercial Emails to contacts mined from PeeringDB will result in a ban and public embarrassment." As of this writing, Cogent's entry in PeeringDB consists simply of its company name, ASN and a note stating "Cogent has been removed from the DB for spamming".⁶⁵

Cogent Communications personnel - including peering coordinators - have avoided attending the most significant network operators' meeting in North America for four years, and they have been systematically excluded from the most important shared database used by peering coordinators. In spite of this absence and exclusion, Cogent continues to survive as a company, since it offers some of the lowest cost bandwidth on the market, and has a substantial customer base. As this case indicates, it is possible for autonomous systems to operate from purely adversarial economic positions, by forcing potential peers to interconnect with them on favorable terms in order to reach their customer base. However, there are limits to such a strategy, especially from the perspective of the Internet as a whole: if all autonomous systems adopted adversarial positions, rather than cooperative and collaborative positions, it is likely that the inter-domain routing would become less stable and less well-connected overall.

The relationships and practices of peering coordinators offer a socially embedded solution to the problem of balancing openness and secrecy in managing the technical/economic relations of interdependence which characterize the inter-domain routing system. In discussing relationships amongst network administrators in the last section, I showed how openness amongst groups of mutually trusting technical personnel is important to stabilizing the inter-domain routing system from a technical perspective. A similar dynamic applies to peering coordinators, from an economic perspective. Openness is required amongst peering coordinators to perform the inter-organizational coordination activity which is integral to their jobs. However, this openness also calls for a kind of closedness, to ensure that sensitive corporate information which is shared with other peering coordinators is limited to the appropriate contexts. To achieve this balance, it is critical for peering coordinators to meet in person, so that they can come to trust one another, and share information verbally, with no written record. The trust relationships required for peering coordinators to function do not exist independent of space, but are rather produced through, and within, the spaces

⁶⁴As detailed in chapter 4.

⁶⁵Cogent Communications' page on PeeringDB is available at https://www.peeringdb.com/private/participant_view.php?id=267, last retrieved Apr 4, 2014. This page is visible to anyone logging in with the user id "guest" and password "guest", as noted on the public login page for PeeringDB.

of peering coordinators' meetings. Paradoxically, the apparent placelessness of the Internet is produced only through activity which must be physically co-located. Trust relationships amongst peering coordinators offer the means through which openness and closedness may be balanced in the practice of peering, to maintain stable relations of interdependence within the inter-domain routing system.

6.6 Operating the Internet

The inter-domain routing system cannot be understood purely in terms of economics and technology. These offer valuable analytical perspectives, but the economics of inter-domain routing cannot be reduced to rational choice, nor can technical issues be addressed purely through technical fixes. As I have shown in this chapter, and the last, the topology of the inter-domain routing system is ordered through a combination of centralized institutional activity, and distributed social relations and practices requiring inter-organizational coordination and collaboration. The technical/economic relations of interdependence amongst autonomous systems in the Internet's inter-domain routing system can only be produced and stabilized through socially embedded solutions.

The sense of embeddedness which I invoke here, in relation to the embeddedness of network interconnection, is twofold. First, as the embeddedness of the market for network interconnection in social relations amongst peering coordinators responsible for negotiating interconnection agreements. This is the the interpretation of embeddedness commonly drawn from Polanyi (2001) and Granovetter (1985). However, there is a second sense of embeddedness which I believe may be drawn from both Polanyi and Granovetter: of the ways in which social relations of trust required for the everyday practice of network administration embed market relations. These are not simply operational relationships functioning over pre-existing economic relationships of network interconnection, but rather must be understood as constitutive of the broader social system of distributed governance which embeds network interconnection markets.

These dynamics are well-recognized within the network administration community, who often frame their activity in terms of "the good of the Internet", implying a notion of the common good of the network administration community, and of Internet users more broadly. This framing carries with it a sense of the Internet as an ideal which must be served by those who produce it; and in doing so, this framing provides a means through which purely economic motivations can be balanced against notions of community, trust and associated activities of coordination and collaboration.

This balance is also produced through pragmatic considerations for peering. Autonomous systems towards the core of the Internet must trust that personnel at their peers are sufficiently "clueful", so that they can "trust" the BGP announcements being relayed to them. Peering coordinators must be able to trust one another in order to be able to share the information required to negotiate peering agreements. Network administrators and peering coordinators need to be able to leverage their relationships in order to resolve issues

which cut across organizational boundaries, or which are concerned with the health of the inter-domain routing system as a whole.

Admission to these communities, and associated networks of trust relations, is based in part upon individual reputation, and in part upon organizational affiliation. In general, autonomous systems which are topologically more central to the inter-domain routing system experience greater risks and uncertainties, and accordingly require personnel who can maintain strong trust relationships with their counterparts in other autonomous systems. Interpersonal social trust relationships are formed to follow the topology of risk and uncertainty surfaced by the technological and economic arrangement of the inter-domain routing system.

Risks and uncertainties which threaten the inter-domain routing system as a whole may be addressed through centralized institutions, which are in themselves often driven by technical personnel. A prime tenet of the RIR system, for instance, is to manage IP address block allocations to minimize the size of the default-free routing table. Similarly, the SIDR WG at the IETF is focused on developing new technical mechanisms to help secure BGP.

The topology of the inter-domain routing system is ordered through distributed relations of trust, which rely on centralized institutions for certain kinds of functions, such as resource allocation and standards development. A variety of other centralized institutions provide other functions which may be opted into on a voluntary basis, such as the Team Cymru bogon feeds. Finally, network operators' groups, such as NANOG, provide the sites for the production and reproduction of relations of trust, and a means to form a sense of community over these relations. Following Cohen (1999), it is this symbolic construction of community which provides a basis over which the ideal of "the good of the Internet" may be constructed: within which volunteer work for the common good may be imagined, whether volunteering to be an office bearer at NANOG or ARIN, or volunteering to help another network administrator in need.

This sense of community is critical to the practice of network administration. The process of learning to be a network administrator, or learning to be a peering coordinator, is intricately tied up with the establishment of relationships within these communities. Even those who have technical backgrounds can only learn the skills and knowledge required for these roles in coordination and collaboration with others in these communities. As I have demonstrated, much of the practice of these communities relies upon coordination and collaboration. Accordingly, the recognition of skills and knowledge - "cluefulness" - is an integral component of the formation of relationships within these communities. These are communities of practice (Lave and Wenger 1991), in which practice is learned and made sense of through the process of becoming a participant in these communities.

It is the combination of these features - trust, community, and practice - which together form the distributed system of governance for the Internet. Centralized institutions - such as the IETF, ICANN and the RIRs - are commonly understood to have political power, and to be sites of contention over power. However, the community of network administrators is rarely understood in this way. I argue that the arrangement of social and economic relations, and technologies, which make up the inter-domain routing system - the community of practice

of network operators - must be understood as a political actor in its own right.⁶⁶

The Internet is commonly framed as a distributed technological system. As I have shown, it is also a distributed social system. It is critical to understand the ways in which power functions within this distributed socio-technical system, the ways in which power is mediated by combinations of social relations and technological artifacts. It is only in coming to such an understanding that we can think more clearly about the governance of the Internet, and the ways in which the design of technologies, and the design of systems of governance, mutually structure one another.

This framing raises two important questions, which I will address in the next chapter. First, how is this distributed system of governance stabilized over space and time? I have dealt with it thus far solely in the North American context, through an examination of NANOG. To understand this system more completely, it is necessary to examine the means through which social relations and practices are extended to other geographies. Second, how does this distributed system of governance behave in relation to stronger economic power, and a more regulatory state? I have examined the internal logic of the distributed system of governance for the Internet. To truly make sense of it, it must also be examined in relation to different kinds of market organization and state power.

⁶⁶This recognition already happens to a degree: NANOG56 featured an Internet governance panel, and NANOG57 featured a presentation from ISOC on the WCIT. In both instances, presenters appealed to the NANOG community for volunteers to get involved in these policy processes as technical experts.

Chapter 7

Distributed Governance III: Ordering Geography

7.1 Disembedding the Internet

When I arrived at my first SANOG (South Asia Network Operators Group) meeting - SANOG18 in Pokhara, Nepal - I found that I was a minor celebrity amongst the organizers. They were surprised that someone from the University of California, Berkeley, would want to attend SANOG; they were certain that SANOG would have no technical knowledge to offer that was new to me. I was met with similar expressions of surprise at the training sessions that I attended, and the main conference sessions. At NANOG meetings, my identity as a student, as an academic, was unsurprising and well-understood. At SANOG meetings, in contrast, I was exotic, even though I am from India, the dominant country in the region that SANOG serves.

In moving from NANOG to SANOG, I had to disembed myself from the more global North American context, and re-embed myself in the more peripheral South Asian context. In the process of doing so, I - and those I encountered at SANOG - had to make sense of my foreignness, and find ways to integrate my identity into common social knowledge. Similar challenges are faced as the elements of distributed governance which I described in previous chapters - relationships, practices, community, and institutions - travel and extend themselves across space. They need to be *disembedded* from the contexts of their origins, imagined as global absolutes, and *re-embedded* in remote contexts.¹ Yet, this process of disembedding and re-embedding occurs over network protocols which *are* global absolutes. The forms of network protocols are necessarily invariant, to allow technological interconnection within and between different contexts. The elements of distributed governance are articulated differently in different contexts; but their articulations are conditioned and constrained

¹I use the term “context” to account for the specifics of a locality: culture, politics, socioeconomic relations, and so on. It is important to note that in my usage, a locality need not be *local* in the sense of physical geography. A context may also be “local” to the infrastructure which provides connectivity at regional and global scales.

by the universal invariance of protocol standards. In this chapter, I examine the work of disembedding and re-embedding of the elements of distributed governance amongst network operators' groups, and the ways in which these elements are altered as they move through this process.²

As I discussed in earlier chapters, the system of distributed governance which orders the topology of the inter-domain routing system operates upon ideals of openness. The institutions of Internet governance are open to participation. Inter-organizational relationships are formed to enable sharing of information, and coordination and collaboration. This ideal of openness operates in tension with senses of closedness: corporations seek to protect their information, and restricted subgroups of the network administrators' community may be formed to share sensitive information. As I will discuss in this chapter, the tension between openness and closedness is complicated further by particular kinds of state involvement and market structure. This results in the articulation of distinctive local governance formations, which must reconcile their distinctiveness with that of the global system of Internet governance. The problem of geographically extending the Internet's distributed system of governance is not limited to disembedding and re-embedding, but must also deal with *integration* across different contexts. Where previous chapters are concerned with the evolution of the Internet across time, this chapter is concerned with the evolution of the Internet across space.

Integration across contexts is required to support the existence of a single global Internet, rather than multiple fractured internets. The risks and uncertainties outlined in chapter 4 are not restricted to any particular context; rather, they are global concerns, carried across time and space in the universal forms of protocol standards. Each of the elements of distributed governance which address these risks and uncertainties require distinct mechanisms for disembedding, re-embedding and integration. Trust relationships require individuals who can travel to become respected members of network operations communities in different geographies. The extension of communities of practice requires the formation of professional organizations which are analogous to NANOG, with linkages to other network operators' groups and the larger system of Internet governance; and the provision of formal training programs, and spaces for informal mentoring. Centralized institutions require formal mechanisms for the extension of hierarchy, to ensure that new regional institutions are in alignment with established institutions. I explore all of these mechanisms of integration in this chapter, through an examination of the Internet in South Asia, focusing in particular on India.

By examining the production of Internet infrastructure in South Asia, I am able to perform a relational comparison with the North American case discussed in prior chapters. This work is relational, in that it offers the opportunity to construct a more complete perspective on the Internet as a global system, by understanding how its production in different geographies is related together to form a global whole. This work is also comparative, since I study variations across these two cases to make sense of the limits and possibilities of distributed

²Giddens (1991) notably uses concepts of disembedding and re-embedding in his theorization of modernity. See Chapter 2 for a discussion of my usage of these concepts, and how my usage differs from Giddens'.

governance, as it manifests in different contexts.

I think of the mechanisms of disembedding, re-embedding and integration together as a system of spatial practice (Lefebvre 1991): the practice which is involved in the production of space, specifically the space of Internet infrastructure spanning North America, South Asia, and the rest of the world. It is through this spatial practice that the virtual space of the Internet comes to be imagined and experienced as independent of physical space, making this spatial practice a critical area of study to understand the politics of the Internet.

In prior chapters I focused on the Internet's technical communities to establish the internal logic of distributed governance. In this chapter, I use the case of India to examine the limits of distributed governance in its interaction with other systems of power, by examining the role of the state, and of market organization, in the spatial practice of the production of Internet infrastructure. I will show how and why the formation of technical community is retarded in India, and then examine the role of the state and market organization in provisioning governance functions which might otherwise be provided through distributed governance.

I attended two SANOG meetings in the course of my research, SANOG18 in Pokhara, Nepal, and SANOG22 in Mumbai, India, which reflect the organization of this chapter. In the next section, I examine the mechanisms involved in the extension of the Internet to South Asia at SANOG, drawing primarily from my fieldwork and interviews at SANOG18. I draw primarily from my fieldwork at SANOG22 and interviews with Indian policy makers. In the following section, in which I examine the production of Internet infrastructure in India through the intersection of state and market concerns with the distributed system of governance of Internet infrastructure. Finally, I conclude with a discussion of the mechanisms through which Internet infrastructure is produced in its peripheries, and the ways in which the Internet's distributed system of governance is re-embedded in peripheral locations in relation to local state interests and market organization.

7.2 Extending the Internet

SANOG is explicitly framed as following in the model of other network operators' groups, such as NANOG, to provide a space for coordination, collaboration and education in the South Asian context. The countries that SANOG serves include Afghanistan, Bangladesh, Bhutan, India, Maldives, Nepal, Pakistan and Sri Lanka.³

South Asia presents an interesting relational contrast to North America. Most of the countries in the region occupy the peripheries of the world system, with the exception of India, which is the dominant power in the region. From a topological perspective, all of the countries in the region, including India, occupy peripheral positions on the Internet, as they are net consumers of content, with comparatively little locally hosted content and applica-

³From the SANOG website, at <http://sanog.org/introduction.htm>, last retrieved Apr 9, 2014.

tions.⁴ The vast majority of autonomous systems in South Asia occupy peripheral topological positions in the Internet's inter-domain routing system. There are a few Tier 1 autonomous systems in South Asia, all of them in India, some of which maintain a presence at NANOG. As I will discuss, the contradictory mixture of peripheral and central positions within the same region has material implications for the social formations of network administration in South Asia.

7.2.1 Organizing SANOG

The first SANOG meeting was held in January 2003, in Kathmandu, Nepal. Meetings are held twice a year, rotating through locations in Bangladesh, Bhutan, India, Nepal, Pakistan, and Sri Lanka. A typical SANOG meeting is split into three segments: two days for the main conference sessions (and associated social events), two days for tutorials, and then a week of technical workshops. As this organization of the meeting program suggests, education is a major focus at SANOG, much like the early NANOG meetings. Unlike the early NANOG, SANOG lacks a large core of senior technical personnel who attend meetings regularly. I perceived attendees at SANOG meetings to be much younger than those at NANOG meetings, reflecting an emerging - rather than established - cadre of technical personnel managing network operations in South Asia.

SANOG operates on a much smaller scale than NANOG, with fewer meetings per year (two, rather than three), and smaller numbers of attendees. It is difficult to estimate exact attendee numbers for SANOG, since publicly available attendee records cover only the main conference, not the tutorials or workshops. Available data indicates that SANOG meetings have had an average of 129 attendees. The highest recorded attendance was at SANOG12 (held in Kathmandu, Nepal, in 2008), which had 220 attendees, and the lowest was at SANOG20 (held in Karachi, Pakistan in 2012), which had 41 attendees.⁵

Attendees may choose which of the three segments of a SANOG meeting they wish to attend, and are charged per segment. In consequence, senior management typically attends the main conference sessions, while more junior technical personnel are sent to the tutorials and workshops. An instructor at SANOG18 commented on the different imperatives of these two groups:

At least what I've noticed is two distinct audiences at SANOG. There's those that come strictly for the workshops, and maybe stay for the conference and

⁴This is in no small part due to the unreliability of local infrastructure and electrical power supply, making it challenging to operate data centers in the region.

⁵These records are themselves incomplete, covering only 17 out of 23 SANOG meetings. Attendance records list only attendees of the main conference, and do not include those registered just for workshops or tutorials. In addition, some regular attendees, who are on the SANOG committee, or acting as instructors for training sessions, are sometimes not listed in these attendance records, as I noticed when comparing attendance details to my field notes. Attendee data, when available, is linked from individual SANOG meeting pages on the SANOG website at <http://www.sanog.org/previous.htm>, last retrieved Apr 15, 2014.

tutorials, and then there's quite a few people who come purely for the conference days, and I guess they're here more to network and all that more than anything. [I27:5]

Attendance fees vary depending upon the venue: at SANOG18 in Pokhara, Nepal, I paid USD 200 for a "passport" (providing access to the conference, tutorials and workshops); a similar package was priced at USD 300 at SANOG22 in Mumbai, India. With hotel and travel costs factored in, those attending SANOG meetings require organizational support to pay their way. The bulk of attendees, who are younger network administrators attending training sessions, may never return to SANOG. Frequent attendees - who are typically on SANOG committees - have substantial experience in the network operations world, often with international linkages, which I will discuss further in the next section.

SANOG is operated by a range of different committees. The core committee is responsible for the overall management of SANOG. The program committee oversees the SANOG conference programs. The fellowship committee manages the allocation of fellowships to attend SANOG conferences. Finally, the advisory committee provides support to all of the other SANOG committees. There are no elections to these committees; individuals are admitted by invitation. Those involved with the SANOG committees span a range of senior technical personnel with roots in the South Asian region (some of whom work in other parts of the world), and others who are well-respected across global network operations communities. The organizations they are affiliated to include local IXPs and ISPs, non-profit entities involved in providing technical expertise and infrastructure (such as PCH and NSRC), and Internet governance organizations (including ISOC and APNIC).⁶

SANOG is supported by a variety of organizations which provide paid sponsorship, technical infrastructure, or instructors for tutorials and workshops. In its early years, SANOG was incubated by PCH, which covered meeting costs and salaries for personnel organizing SANOG [EMAIL:06-28-2014].⁷ Over time, the SANOG funding base has become more diverse. Local ISPs and ISP associations often provide financial support, as do other private entities such as Google, Cisco Systems and Juniper Networks. Internet governance organizations such as ICANN, ISOC and APNIC provide financial support, as well as instructors. Non-profit entities, such as PCH and NSRC, often provide instructors for training sessions. The SANOG website and email list are hosted by ISC, which is also a non-profit entity. Local governments sometimes provide support: SANOG22 was co-hosted the ISP Association of India (ISPAI) and the National Internet Exchange of India (NIXI), and counted the Maharashtra state government as a supporter. Mumbai - where SANOG22 was held - is the capital of the Indian state of Maharashtra; NIXI is operated as a public-private partnership between the Indian government and Indian ISPs. In fact, it was announced at SANOG22

⁶For a list of committees and committee members, see <http://www.sanog.org/committee.htm>, last retrieved Apr 18, 2014.

⁷As part of its mission in serving the Internet's technical communities, PCH has supported the development of several other network operators groups, including CaribNOG (Caribbean Network Operators Group) and MENOG (Middle East Network Operators Group) [EMAIL:06-28-2014].

that NIXI and ISPAI would henceforth host one SANOG meeting in India every year, out of the two meetings per year.⁸

As this account suggests, SANOG presents a local nexus of the system of relationships amongst organizations involved in the production of the Internet in a variety of scales and functions. SANOG is made possible only by imagining it as being similar in form to NANOG and other network operators' groups. The instructors who provide training at SANOG have roots in, and connections to, other regional and global contexts even if they are from South Asia, and are typically supported by global Internet governance organizations and non-profits. Financial support for SANOG is provided by global entities, as well as government agencies and the local Internet industry. NANOG presents a similar nexus of relationships, but is much more independent of other organizations in this system: NANOG is a peer to other organizations involved in producing the Internet, rather than reliant on them for its existence. This position changed through NANOG's history, as it went from being government-supported (through NSF funding to Merit) to becoming independent from Merit (with a loan from ARIN).

The differences between SANOG and NANOG can in part be explained as a matter of organizational maturity, but cannot be reduced to this. Rather, these differences represent differently positioned state interests, and functional and topological variations. I will develop these themes in the remainder of this chapter, beginning with an examination of the ways in which interpersonal relationships and community are established at SANOG, within South Asia, and between South Asia and the rest of the world.

7.2.2 Extending Relationships

As I noted earlier, SANOG was incubated by PCH in its formative years, with support from a small group of network administrators in South Asia who perceived a need for network operations training in South Asia, and had already been involved with such activities in the South Asian region. One of those involved with the founding of SANOG, who was employed by PCH to manage early SANOG meetings, told me how his involvement with network administration, training and network interconnection in Nepal led him to an interest in founding SANOG. He taught himself networking in college as he realized that there was power associated with being a network administrator, in the ability to control and provision access to the systems he managed. He went on to be involved with operations for several networks and ISPs in Nepal. As a result of his experience, he was hired as an instructor for a program administered by Cisco and the United Nations Development Programme (UNDP) aimed at providing computer networking education across colleges in Nepal. This led to connections which allowed him to attend the INET conference in Stockholm in 2001, and then the Asia-Pacific Regional Internet Conference on Operational Technologies (APRICOT)

⁸This arrangement has also been announced on the SANOG website, at <http://www.sanog.org/hosting.htm>, last retrieved Apr 17, 2014.

in 2002.⁹ Thanks to the connections and knowledge he gained from these international fora, and realizing that traffic between autonomous systems in Nepal was routed through expensive international routes, he went on to set up the Nepal Internet Exchange (NPIX) with support from PCH. He currently works in Singapore [I26:1-3].

This story is representative of the experiences of many of the senior technical personnel who attend SANOG meetings. As individuals, they established regional and international relationships through employment, interactions with non-profit entities with a global presence, and attendance at regional and international technical conferences. In many ways, these relationships are not vastly dissimilar from those I described in the context of NANOG. These senior technical personnel leverage these relationships in their everyday practice, to maintain local infrastructure (private to their organizations, and shared at local IXPs), to engage in regional and global governance activity (at APNIC and APRICOT), and to create and sustain SANOG. Some of these individuals continue to reside and work in their home countries, while others have found employment outside South Asia. The network of relationships amongst organizations which provide support to SANOG exists only through the interpersonal relationships established by these key senior technical personnel, spanning different scales of regional (across South Asia and the Asia-Pacific regions) and global geographies.

This core of senior technical personnel is clearly identifiable at SANOG meetings. The core senior technical personnel form a distinct connective group, constantly interacting with one another, sometimes casually, and sometimes in formal settings, such as committee meetings or conference panels. In my conversations with attendees from Bangladesh, Bhutan, Nepal and Pakistan, it became clear that they looked up to these senior technical personnel, speaking of them with reverence.¹⁰ In every instance, multiple attendees from the same country identified one or two senior personnel by name who they saw as being key to the development of the Internet - and attendant technical community - in their country. In social settings - such as the lobby of the conference venue - I often found the core of senior technical personnel surrounded by attendees from their home countries. The members of the core group are anchors for the network operations community in their countries, just as much as they provide social connectivity to regional and international contexts.

Membership in the core group is formed and renewed through social occasions, as much as through technical competence, just as it is at NANOG. A SANOG committee member from Bangladesh praised the hospitality of his Nepali hosts at SANOG18:

I am having parties every night, damn good parties. With whom? Only with my Nepali friends. I have a bunch of 12, 13 people here every night. They collect me from anywhere and take me to new places. . . . they're making fun of different

⁹The INET conference series is run by ISOC, and covers issues of technology and policy. APRICOT is analogous to NANOG, covering business and technical issues related to Internet infrastructure across the Asia-Pacific region.

¹⁰There were no attendees from Sri Lanka at either of the SANOG meetings I attended. I encountered one attendee from Afghanistan at SANOG22, but his responsibilities were limited to managing a college campus network. I will restrict my comments on Pakistan, since attendance from Pakistan was extremely limited, although I did conduct an interview with a senior network administrator from Pakistan.

things and making me laugh and enjoy the moment, everyone. They will come to Bangladesh [and I will entertain them in turn]. If I go to Bhutan, these people will come, and whoever took part in the class, the attendees, they'll come. We actually have a very good community here, except for India. [I29:4-5]

Similarly, a senior network operator from Bhutan told me how he and other members of the SANOG core group went out drinking together every night at various SANOG meetings, even chartering a bus to travel together from Kathmandu to Pokhara for SANOG18 [F-SANOG18:142]. Hospitality offered by those from a country hosting a SANOG meeting is reciprocated as they travel to other countries in the region. The social obligations incurred and fulfilled through this process are constitutive of the thick relationships formed amongst the core group at SANOG. This interviewee, and others, spoke of how India is not well-represented at SANOG. While there are senior technical personnel from India who play important roles at SANOG, there is less of an awareness of their prominence within India, although they are well-regarded by their peers at SANOG. For instance, the SANOG co-founder I interviewed told me how he learned the basics of inter-domain routing at a training session conducted by an Indian trainer, who was then employed with Cisco Systems, and is currently a SANOG committee member [I26:2].

The variations in these individuals' social positions in relation to other technical personnel from their countries can be explained by the different roles that these individuals play in establishing interconnections between autonomous systems in their home countries, and between their home countries and the rest of the world. In India, these functions are carried out by distinctive arrangements of government and private industry. In contrast, topologically peripheral countries such as Bangladesh, Bhutan, Nepal and Pakistan rely on senior technical personnel from their countries to provide leadership in setting up IXPs, and facilitating international connectivity.¹¹

A second tier to the core of attendees is formed by technical personnel who are responsible for managing significant autonomous systems in their countries, or involved in IXP operations. Many of these attendees have limited international relationships of their own, in the process of establishing international connectivity for their autonomous systems. Others are being groomed to take on leadership roles in their countries by the core senior technical personnel at SANOG. This second tier of attendees is gradually drawn into the orbit of the core group, often through invitations to meals or drinks, similar to the manner in which relationships are established at NANOG. For instance, a Bhutanese network administrator whose acquaintance I made at SANOG18 was frequently called away by his manager to meetings with members of the core group.

Outside of the core group of senior technical personnel, most SANOG attendees self-segregate, first by organizational affiliation, and then by country of origin, with limited interaction with attendees from other South Asian countries. From a functional perspective,

¹¹There are, of course, substantial variations in the ways in which these facilities are constructed in different countries in South Asia, which I will discuss only in a limited manner. I will explore the nature of these arrangements in detail for India in the next section.

attendees from different organizations who do not need to interact with one another as part of their everyday practice have little incentive to establish cross-organizational relationships. They rely on senior technical personnel from their countries to perform inter-organizational bridging work when required, such as for the formation and maintenance of IXPs, and country-level ISP industry associations. That said, they do form relationships with others from their country thanks to SANOG meetings, as one of the core group of senior technical personnel pointed out:

Like a lot of the guys, they all work for competitors in the country. Whenever they meet outside of that confined space, generally, they never used to meet each other. Now they get to meet and they have a direct line of communication, and they talk a lot more with each other. . . .they might not know each other at all back home. Here, they are all friends. [I26:4]

These relationships do not extend to the SANOG mail list, on which traffic is scant. The SANOG list has had, on an average, just 19 messages per month since its inception in 2003. The majority of these messages are notices, such as security advisories or invitations to SANOG and APRICOT conferences, with very little by way of the extended discussions visible on the NANOG mail list.¹² The SANOG mail list is not a site through which South Asian network operators maintain their relationships and community, reflecting the sparseness and intermittent nature of relationships and community in South Asia outside of SANOG meetings. The formation and maintenance of relationships at SANOG depends very much upon attendees being co-located in the shared physical space of the conference, even more so than at NANOG.

Segregation of attendees is also a consequence of attendee numbers being dominated by the host country, with substantial variations in the number of attendees from other countries. It is expensive to travel within South Asia, and it can be difficult to obtain visas between certain countries. SANOG meetings are intentionally held across different countries in South Asia to allow for greater participation from the host country. Of the 23 SANOG meetings which have been held to date, more meetings have been held in India than have been held in the other South Asian countries: 6 meetings in India, against 3 or 4 in each of the other countries. Table 7.1 provides an analysis of country-level participation at different SANOG meetings for which attendance data is available.

Although the data from which this table is constructed is limited in several ways, it is nonetheless indicative of patterns of participation across the region.¹³ At the most basic level of analysis, it is clear that most attendees are drawn from the host country. This number is lowest for Bhutan, which is an outlier as being considerably smaller than any other country in the region, with a proportionally smaller population of network administration

¹²Archives of the SANOG mail list are available at <http://blog.gmane.org/gmane.org.operators.sanog>, last retrieved Apr 21, 2014.

¹³See footnote 5 for details of the raw data behind this analysis. Attendance from Afghanistan and Maldives - which are nominally in the region serviced by SANOG - is extremely limited, so I have included them in the “Other” category.

	Bangladesh	Bhutan	India	Nepal	Pakistan	Sri Lanka	Other
Bangladesh 506 attendees over 4 meetings	87% 441	1% 4	1% 6	5% 25	1% 7	1% 5	4% 18
Bhutan 209 attendees over 2 meetings	24% 50	43% 90	5% 11	11% 22	1% 2	0% 0	16% 34
India 485 attendees over 3 meetings	9% 45	1% 7	69% 335	7% 33	1% 5	2% 8	11% 52
Nepal 479 attendees over 3 meetings	14% 65	4% 18	6% 31	63% 301	2% 8	2% 8	10% 48
Pakistan 292 attendees over 3 meetings	6% 17	1% 3	0% 0	5% 14	79% 230	1% 3	9% 25
Sri Lanka 228 attendees over 2 meetings	8% 19	2% 5	4% 8	7% 16	6% 13	66% 150	7% 17

Table 7.1: The proportions of attendees from different countries attending SANOG meetings in a given host country. The rows are meeting host countries, and the columns indicate the country of origin of attendees. Cells have the percentage and total number of attendees from a given country across all meetings for which records are available in a host country.

personnel. Attendees from countries in the “Other” category are typically instructors for training sessions, traveling from Europe and North America, often sponsored by PCH or NSRC, or representatives of APNIC, traveling from their offices in Australia.

More interesting are the numbers which indicate a willingness to take on the costs and difficulties of traveling to other South Asian countries to attend SANOG meetings. Attendees from Bangladesh show the greatest willingness to travel, with strong ties to Bhutan, India and Nepal. Attendees from Nepal are the next most willing to travel, with strong ties to Bangladesh, Bhutan and India. Attendees from Bhutan show a willingness to travel to Nepal, but not so much to other South Asian countries. Attendees from India show some willingness to travel to Nepal, but not so much to other South Asian countries, and not at all to Pakistan. This is likely because on the one hand, Indians require no visa to enter Nepal, and on the other hand it is difficult for Indians to obtain visas to enter Pakistan. Attendees from Pakistan and Sri Lanka are reluctant to travel to other South Asian countries for SANOG meetings, probably due to considerations of cost (especially in the case of Sri Lanka) and difficulties associated with obtaining visas.

This data shows strong ties bilateral ties between Bangladesh and Nepal, and between

Bhutan and Nepal, as indicated by the numbers of attendees who are able to have their organizations sponsor the costs of travel and the effort of getting a visa between these countries. This grouping of countries reflects a cluster which was apparent from my observations and interviews at SANOG meetings. When I asked SANOG committee members about this pattern, they included Pakistan in this cluster, saying that this cluster was formed by a similar set of concerns faced by peripheral countries in South Asia. These concerns are primarily related to issues of international connectivity. Bangladesh and Pakistan are able to use submarine cables for their international connectivity, since they have coastlines at which cables can land. In contrast, Bhutan and Nepal have to depend on large Indian ISPs to provide them with optical fiber links to submarine cables landing in India.¹⁴ The physical topology of the Indian subcontinent, defined by territorial boundaries and coastline, has a strong influence on the economic arrangements for interconnection - and resulting topological positions - of autonomous systems in the inter-domain routing system.

The peripherality of these countries is not merely a matter of physical location; these are all economies which are less developed than India, with an accompanying perception of a lower level of development of a sufficient numbers of network administrators with adequate knowledge and skills. As a Bangladeshi interviewee, who is on the SANOG committee, commented:

... in India, they are much more developed in the technical sense than us. ... They get chances to train themselves, but it is not at SANOG. Maybe it's in the USA. Maybe it's in Singapore, or some other bigger places; they go to APRICOT. ... if SANOG is in India, we'll see lots of participants [from India], but other than that [very few from India].

India has four large backbone autonomous systems, all of which also provide consumer Internet access. Of these, one is a government-run telecommunications provider. The other three are privately operated, maintaining national and international infrastructure. None of these dominant South Asian autonomous systems has been a significant presence at SANOG. Out of 2200 attendees across all SANOG meetings for which attendee records are available, just 42 attendees are from these four autonomous systems. Of those 42 attendees, only 2 attended a SANOG meeting held outside India. The bulk of Indian representation at SANOG is from smaller autonomous systems, who themselves occupy the peripheries of the topology of the inter-domain routing system in India.¹⁵ Rather than attending SANOG, the large Indian autonomous systems attend fora at which they may establish the relationships they need to operate at a global level, such as NANOG and APRICOT. This point was brought home to me at SANOG18 in Nepal, when a senior Indian network administrator,

¹⁴International connectivity via satellite is also a possibility, but is not a realistic option for high speed Internet connectivity, which is why I do not discuss it here. All countries in South Asia obtained Internet connectivity almost entirely via satellite at one point or another in their history, but this has changed to favor lower cost optical fiber connections.

¹⁵Some of the Indian representation is also from services and consultancy companies which provide outsourced manpower to operate network infrastructure.

who is on the SANOG committee, purposefully sought me out, commenting “we don’t see many Indians at SANOG” [F-SANOG18:108].

My insider/outsider status made me a bridge between attendees from different countries. This was partly because of my perceived exoticism - as a doctoral student at the University of California, Berkeley - which drew several attendees from different countries who wanted to get to know me better, and understand what it was like to study in the USA. This was also because I quite consciously sought out attendees from different countries for research purposes. By the end of my first SANOG meeting, I frequently found myself in the company of attendees from Nepal, Bangladesh and Bhutan, as well as trainers from APNIC, and members of the core group of senior technical personnel. In comparison, my presence at NANOG had little effect on social connectivity; rather, I had to find my way into the social world of NANOG.

While the effect of my presence at SANOG was certainly limited, the differences between the nature of my engagements in SANOG and NANOG are indicative of the differences in the structure and strength of relationships in these distinct technical communities. Where NANOG is patterned with a dense web of thick relationships amongst a sizable core of attendees, SANOG depends upon a small number of individuals to act as bridges between organizations and countries. I contend that this is largely due to the difference in the topological arrangement of autonomous systems represented at SANOG and NANOG. Relationships formed at NANOG often have immediate implications for everyday practice, due to the denser interconnectivity of autonomous systems in North America, and greater risk and uncertainty in inter-domain routing, which call for stronger trust relationships. In comparison, autonomous systems represented at SANOG rely on large autonomous systems in India, or the Asia-Pacific region, to provide them with international connectivity. These are primarily transit, rather than settlement-free, peering relationships, in which risk and uncertainty may be mitigated by the use of BGP filters. Many South Asian autonomous systems do engage in settlement-free peering relationships at country-level IXPs, but the amount of traffic exchanged at these IXPs is relatively low compared to that which flows over international links, resulting in lower risk and uncertainty in these peering relationships.

In consequence, trust is of less importance at SANOG than it is at NANOG, for the relationships required in the everyday practice of network administration, since these relationships are involved in managing customer-provider, rather peer-peer, relations between autonomous systems. The relationships formed at SANOG are salient for the creation and operation of IXPs, which require business and technical personnel at competing ISPs to come to terms and cooperate with one another. Yet these relationships are themselves made possible by leadership from the limited number of individuals who occupy the core group of senior technical personnel at SANOG. While these relationships can, and do, grow towards being interpersonal trust relationships, they are at first formed through cooperation which is anchored by a few respected senior technical personnel.

The attendance at SANOG meetings represents the peripheries of the Internet in South Asia, whether as autonomous systems in peripheral countries, or as peripheral autonomous systems within the central South Asian country, India. SANOG meetings provide a common

physical space at which to address the common concerns of those who populate the periphery of the inter-domain routing system in South Asia. As such, SANOG meetings represent a spatial construction which spans the territorial spaces of the countries of South Asia, and inverts the hierarchy of topological positions of autonomous systems. Where NANOG meetings are mostly populated with Tier 1 and Tier 2 autonomous systems, SANOG meetings are mostly populated with peripheral autonomous systems. Although SANOG is imagined as being analogous to NANOG and other network operators' groups, the disembedding, re-embedding and integration of relationships at SANOG results in a distinctive formation of peripheral participants, which mirrors the political geography, and topological arrangements of inter-domain routing, in South Asia.

7.2.3 Extending Practice

If trust is less salient to the formation of relationships at SANOG meetings, what is the content of these relationships, and what work do these relationships do? Rather than being formed for the purpose of engaging in the everyday practice of network administration, relationships at SANOG meetings are primarily formed in the process of, and for the purpose of, *learning* the practice of network administration. A Bangladeshi instructor told me how he perceived these ties to be formed in the process of conducting training sessions at SANOG:

After becoming an instructor, I saw most of the students sending me mails with different queries. When I responded, they were encouraged and asked me more questions. I have limitations [to my knowledge], you know, so I passed these questions to certain guys who are experts. What happened when they [experts] responded, the students were so encouraged; so there is a tie, an inter-regional tie, among the instructors, speakers, attendees. [I29:4]

There is a hierarchy to the relationships at SANOG, providing a connection between the most peripheral, least experienced attendees, and experts from outside the region, channeled through instructors. Without exception, instructors are experienced technical personnel, from South Asia and from other parts of the world, who participate in the system of trust relationships linking South Asia to the rest of the world. Those who are from South Asia are part of the core of senior technical personnel involved in organizing SANOG and providing technical leadership within their home countries. The relationships that junior technical personnel form in the course of learning their craft through SANOG cannot be characterized as trust relationships. However, these relationships do form the basis for engaging in the practice of network administration, and building a stronger imagination of being part of a global community of network administrators. Over time, depending on the professional trajectories of junior technical personnel attending SANOG, the relationships formed in the process of learning the practice of network administration may form the basis for thicker relationships which include the degree of trust required amongst network operators involved in managing more central autonomous systems, or even for the formation and maintenance of regional interconnectivity at local IXPs.

Technical training sessions occupy the majority of the time at SANOG meetings, including week-long workshops, and two days of tutorials. This focus on training reflects the peripheral position of most SANOG attendees, and their associated need for training.¹⁶ The workshops and tutorials address a variety of subjects requiring different skill levels, with multiple tracks running in parallel. Subject areas frequently shift from one SANOG meeting to another, including inter-domain and intra-domain routing, network security, DNS operations and more. One of the founders of SANOG - who has taught many routing workshops - commented on the importance of training at SANOG:

The first time I taught this [routing] workshop in Bangladesh, I had 51 people in the room. I was teaching it alone. Then gradually you get more and more people. . . . the routing workshop seems to be more and more popular, no matter where and when. It says basically that you're getting more newer people coming in, more and more new participants entering the system. Then people generally gradually are going to more advanced workshops like MPLS or multicast or something like that.¹⁷ Yeah, so it actually probably works a lot better in South Asia than even in the U.S., or even at the APRICOT meetings, because there is no such training there. [I26:3-4]

Although the topics covered by tutorials and workshops have varied from one SANOG meeting to another, a workshop concerned with routing has always been offered, covering both exterior gateway protocols (which is to say, BGP) and interior gateway protocols (IGP) such as OSPF.¹⁸ As the popularity of the routing workshop indicates, routing is perceived to be the basic skill that network administrators need to perform their work. Once a degree of proficiency is acquired with routing, SANOG attendees may go on to study more advanced topics, which are typically concerned with routing (such as MPLS and multicast), network security, and the provision of services (such as email and DNS) in an ISP. As my interviewee notes, the continued popularity of the routing workshop reflects an ongoing and growing demand in South Asia for the skills and knowledge required for the practice of network administration.

I attended the routing workshop at SANOG18. Two other workshops were held in parallel, covering campus network design and MPLS.¹⁹ The routing workshop was collaboratively taught by two instructors, from Nepal and New Zealand. The Nepali instructor was on the

¹⁶In comparison, while NANOG meetings do offer training sessions, these offerings have always been limited relative to the main conference sessions.

¹⁷MPLS stands for Multi-Protocol Label Switching, a technology used to direct data traffic along labeled paths, rather than by lookups to routing tables.

¹⁸The earliest available website for a SANOG meeting - SANOG3, held in 2004 in Bangalore, India - indicates that workshops on DNS and BGP multihoming were offered. See <http://www.sanog.org/sanog3/workshopapp.txt>, last retrieved Apr 28, 2014. Subsequent routing workshops have covered both BGP and various IGPs. OSPF (Open Shortest Path First) is amongst the most popular routing protocols used within (rather than between) autonomous systems.

¹⁹See the SANOG18 program, available at <http://www.sanog.org/sanog18/program.htm>, last retrieved Apr 28, 2014.

SANOG committee, while the instructor from New Zealand was an employee of PCH. Their training materials were largely adapted from a format created by a former Cisco Systems employee who was instrumental in setting up ISP and IXP workshops for Cisco Systems in the 1990s.²⁰ The instructors were supported by three experienced network administrators - two from Nepal, and one from Bangladesh - who were on hand to help students work through their problems.

The students at the workshop were, almost without exception, male and young. Most students appeared to be in their early 20s, with a scattering of older people. There were just three women in a room of about 50 students. In my conversations during the week of the workshop, it became clear that the majority of my fellow students had been sent to SANOG18 expressly for the purpose of attending this workshop; they would not stay on to attend the tutorials or the main conference [F-SANOG18:6]. Although I was not in a position to establish the numbers of attendees from different countries, it appeared that the majority of students were from Nepal, which was unsurprising, considering that SANOG18 was held in Nepal. The next largest contingent was from Bangladesh, followed by a fair representation from Bhutan, and a small number of Indians.

Many of my fellow students fell into the profession of network administration by accident, as did many attendees at NANOG. A network administrator from Bhutan, for instance, told me how he trained as an electrical engineer, intending to return to Bhutan to work on hydroelectric projects.²¹ After he got his degree, he interviewed with the hardware side of Bhutan Telecom²², and then stayed and moved on to the ISP business [F-SANOG18:46]. A Nepali network administrator I spoke with told me how he used to teach computers to high school students, running a small network for the school, before moving on to work on digital TV [F-SANOG18:91]. Like the SANOG founder I spoke with, many of those attending the routing workshop became network administrators almost without meaning to, and always in response to practical concerns.

The workshop was intended to teach IPv4 and IPv6 routing, using BGP and OSPF. It was constructed as a series of short lectures, with the majority of the week devoted to hands-on configuration and troubleshooting of routing infrastructure. The format of the workshop was intentionally constructed to require collaboration. We were seated at two long tables, each representing an independent routing infrastructure, and split into teams of two. Each team was required to work together to manage a router which was part of the network infrastructure for the table.²³ It was fortunate for me that we were paired for the workshop, as some prior knowledge of the command line interface used by Cisco routers was

²⁰These materials are in wide use in workshops held across the world: many NANOG attendees (some of whom have been trainers themselves) praised these materials as “the gold standard” for training whenever I mentioned them.

²¹India provides investment in Bhutan to construct hydroelectric dams; Bhutan in turn exports electric generated from these dams to India.

²²The government-run telecommunications provider in Bhutan.

²³These were not physical routers, but rather software routers running in virtual machines on computers brought by the instructors.

assumed. Although I have a theoretical knowledge of routing, I had no practical experience with operating routers. I learned how to route with help from my teammate and others at my table, as well as by working through the exercises. The instructors' presentations largely dealt with theory, and not so much with the mechanics of configuring routers.

After the initial setup for the router, we were required to configure interconnections with other routers to create a variety of network topologies using OSPF to set up both IPv4 and IPv6 connectivity. As the workshop progressed, our exercises became more complex, as we were required to group our routers into separate autonomous systems, and then use BGP to interconnect our autonomous systems. In the process, we learned a variety of accepted best practices, including approaches to constructing addressing schemes for networks, reliable methods for configuring OSPF and BGP, and ways to secure networks. In order to complete our exercises, we had to work with our teammates. As the exercises became more complex, we had to walk around our tables to work with teams managing neighboring routers and neighboring autonomous systems to coordinate our configurations to ensure proper interconnections between our routers and autonomous systems. More often than not, this involved testing connectivity from both ends of a link, and then working with other teams to troubleshoot our configurations. In many ways, the format of the workshop created a microcosm of the relationships, coordination and collaboration required for real network operations. The instructors stressed the collaborative nature of the Internet in their presentations. For instance, they pointed out that network interfaces which carry egress traffic can be controlled by filtering routing announcements; in contrast, ingress traffic can only be controlled through cooperation with upstream network providers [F-SANOG18:22].

During the exercises, the instructors were available to answer questions as we ran into problems. After each exercise was complete, the instructors checked our routing configurations in front of the whole class, talking us through the process of testing and troubleshooting the routing infrastructure that each table had set up. The command line interface they used was projected at the front of the room (see figure 7.1), so that we could observe the mechanics of the commands they issued to elicit details of network configuration and fix it when they found it broken, while they talked about what they were doing, and why they were doing it.

The workshop was constructed as a form of apprenticeship: we learned the practice of network administration by doing network administration, and then having our mistakes corrected by experienced network administrators, who illustrated how they thought through this process. As one of the instructors told us during the last day of the workshop:

We're different from Cisco or Juniper instructors, everything we're talking about is based on running networks, rather than describing router features. [F-SANOG18:56]

As at NANOG, there was a stress on the non-commercial aspect of this training, an insistence that it was focused on practical experiences from actually operating networks, rather than teaching attendees how to configure products from particular vendors. A senior network operator from Bhutan echoed this perspective, telling me how he perceived the training at SANOG to be superior to training from vendors in its relevance to practical



Figure 7.1: At the SANOG18 routing workshop.

network operations, especially since instructors are willing to stay in touch with students even after the training is completed [F-SANOG18:142]. The practice of “running networks” cannot be reduced just to technical skills; as the instructors gave their presentations, they also told stories of the practical world of network administration.

The instructors spoke of BGP community tags and MED (Multi-Exit Discriminator) attributes as examples of trust relationships, since these are effective at controlling ingress traffic only if upstream network providers honor the intention behind these attributes [F-SANOG18:60].²⁴ They also made specific comments about the observed behaviors of ISPs in South Asia, pointing out that South Asian ISPs are known for fragmenting their IP address prefixes to the maximum extent possible, announcing thousands of prefixes where just one

²⁴BGP community tags are used for a variety of purposes, such as indicating to an upstream network provider that a BGP announcement should not be propagated to its peers. The MED attribute is used to indicate a preference to upstream network providers about how to select between multiple potential ingress paths.

or two would do. They used such stories to reinforce notions of best practice, commenting that even if an autonomous system needs to announce fragmented address space, it can also announce a covering prefix (aggregating all the fragments), and tag the fragment announcements with the BGP “no-export” community, to prevent them from being propagated by the upstream network provider [F-SANOG18:62]. The importance of interpersonal trust relationships also came up in these stories: during a tutorial at SANOG22, an American instructor stressed the need for building “trust groups” with other network administrators across the world to help resolve issues in other autonomous systems, and other countries [F-SANOG22:465].

Instructors use these stories to establish ideas of the norms and expectations which are embedded in the relationships and practice of network administration. The cooperation and collaboration required to manage inter-domain routing can only work insofar as the majority of network administrators conform to accepted best practices learned from the practical experience of operating autonomous systems. The routing workshop was as much a means to learning the skills and knowledge required for routing, as it was a means to communicate the established norms and expectations which help stabilize the inter-domain routing system.

The workshop was by no means a friction-free experience in the transfer of knowledge and skills. Some students expected a different focus for the workshop than the one adopted by the instructors. In my conversations, several students complained that there was too much theory; they didn’t see why the instructors’ presentations needed to be as long and involved as they were. A student who was a network administrator at a university campus in Nepal was frustrated by the amount of time spent on BGP; he would have much preferred a course which focused entirely on different IGP configurations, which were of most interest to him for his work. A student at a large ISP in Bangladesh - who was amongst the most technically proficient students at the workshop - came expecting a greater focus on IPv6. As with any workshop, students had a range of experience levels and interests, which could not all be treated equally.

The complaints about too much theory were in part driven by issues of language and accent. All of the students spoke English only as a second language, and several of them had limited proficiency in spoken English. Many of these students might have found it easier to spend more time reading and understanding the written materials, than trying to make sense of the instructors’ presentations. At one point, students were having difficulty understanding the accent of the instructor from New Zealand, leading the instructor from Nepal to take over the presentation. I observed similar difficulties of language at SANOG22, when network administrators from a small ISP in North India tried to ask an instructor from South India for advice. The South Indian instructor spoke English fluently, but very little Hindi, while the North Indian network administrators spoke Hindi, but very little English. Differences of language can make it more difficult to establish relationships between South Asia and the rest of the world, and even within South Asia.

The relationships formed at SANOG are perhaps better characterized as a network of practice (Brown and Duguid 2000), rather than a community of practice (Lave and Wenger

1991). The relationships formed in learning the practice of network administration at SANOG are thinner than the relationships formed at NANOG. Junior personnel receiving training do have the opportunity to become more central in the system of relationships at SANOG, but likely will only be able to do so by rising to senior positions within their organizations, and by moving to autonomous systems which are more central to the Internet within their countries or outside South Asia. SANOG represents a nexus of relationships which enable legitimate peripheral participation (Lave and Wenger 1991) in the practice of inter-domain routing. However, peripheral participants in South Asia cannot grow to more central positions purely through gaining skills and reputation, but rather must also shift their organizational affiliation to autonomous systems in more central positions in the inter-domain routing system, thereby acting as conduits between South Asia and the global Internet.

It is possible that SANOG will grow to mirror NANOG over time, forming a stronger community of practice; but this is likely to happen only if the topological arrangements of autonomous systems in South Asia become more densely interconnected, calling a larger cadre of senior technical personnel who form strong trust relationships with each other, and with similar technical personnel outside South Asia. If this possible future were to come to pass, it would paradoxically lead to an inversion of SANOG: rather than providing a space for training and collaboration amongst peripheral autonomous systems, SANOG would become a site for coordination and collaboration amongst more central autonomous systems. In many ways, SANOG is exactly what it needs to be today, representing the needs and interests of the periphery of the Internet in South Asia, and providing institutionalized access to channels between peripheral autonomous systems in South Asia and the global Internet.

The learning of practice at SANOG is a means through which relationships are formed amongst SANOG attendees, and between SANOG attendees and network administrators outside South Asia. However, the varied cultural contexts of South Asia, and the peripheral positions of most SANOG attendees to the inter-domain routing system, make it difficult to characterize the practice, and the relationships formed in learning practice, as being constitutive of community.

7.2.4 Extending Community

If a sense of community is not formed through the shared practice of network administration, and the mutual dependencies and relationships formed as part of this practice (as it is at NANOG), in what ways do network administrators in South Asia imagine themselves as being part of a community, both regional and global? I argue that community is constructed at SANOG primarily through shared ideals and symbolic resources, which are integral to the everyday practice of network administration, and shared concerns, which characterize the peripheral position of autonomous systems in South Asia.

At a basic level, network administrators around the world share a common technical language. Anyone engaged in the practice of network administration must know what an

IP address is, must understand basic concepts of routing, and must be aware of routing protocols such as BGP and OSPF. In the process of developing skills and knowledge in these areas - through training at SANOG, for instance - network administrators also learn how to talk about these technologies with one another. During the workshop I attended at SANOG18, the instructors taught the technology of BGP route filters, and also talked about why filtering was required, and in what situations it is used, explicitly mentioning BCP38, the “best current practice” IETF document which details the use of, and need for, route filters. To understand the technology of configuring a filter, then, is to understand why the technology is required, providing a particular frame through which to comprehend and talk about inter-domain routing. Amongst IETF documents, BCP38 is perhaps the best known in the world of network administration. It is amongst the ideas that are assumed as basic knowledge in conversations at NANOG; when BCP38 is mentioned, it is assumed that others in the conversation are aware of its import, and its meaning in their practice. In bringing BCP38 into the routing workshop at SANOG, the instructors were doing more than just teaching a “best current practice”; they were helping enlarge the conceptual vocabulary with which students think about their practice.

To perform the practice of network administration, a common set of tools are required, which have embedded within them ideals of openness. Tools such as RouteViews, RIPE Atlas and Looking Glass²⁵ depend on the willingness of network administrators around the world to share data about their autonomous systems, and in doing so, make it possible for themselves and others to more easily diagnose the problems they encounter in the practice of network administration. During a tutorial on Internet measurement tools at SANOG22, the presenter drew the connection between tools and the ideal of openness explicitly, pointing out that “publicly available resources make the Internet a nicer place to do your job”. He strongly recommended peering with RouteViews, hosting network probes such as the RIPE Atlas, and maintaining up-to-date data in PeeringDB [F-SANOG22:197-201]. These tools are at once a means to an end, for the practice of network administration, and shared symbolic resources embedding ideals of openness, which allow network administrators to imagine themselves more concretely as part of a global whole.

Ideas of what the Internet is, and how it should be governed run throughout SANOG. These ideals are explicitly espoused during training sessions and presentations, and shape the nature of interactions at SANOG. One of the instructors opened the SANOG18 routing workshop by saying:

Who runs the Internet? No-one. Definitely not ICANN, the RIRs, or the USA. . . . By and large everything evolves for the common good. . . . These sorts of meetings [SANOG] keep things working. At the end of the day, engineers have to support one another. [F-SANOG18:12-17]

The idea of the Internet as a globally distributed system - both in terms of technology and governance - travels to SANOG through statements such as this one. The instructor

²⁵See Chapter 6 for a discussion of these tools.

argued that there is no center of power for the Internet; rather, the Internet works by virtue of “these sorts of meetings” - SANOG, NANOG and other network operators groups - which allow engineers to coordinate and collaborate to “keep things working”. These kinds of statements, which are part of common parlance amongst network administrators, establish the image of SANOG as being equivalent to NANOG and other network operators groups, at the same time privileging the position of network operators groups in relation to countries, corporations and centralized governance institutions.

By exposing students to these shared symbolic resources and ideals, instructors provide the means for students to think about their practice as part of a global system. Through the deployment of these shared symbolic resources and ideals in their everyday practice, students are able to show how they are part of a global system of practice. In doing so, they become able to more meaningfully engage with more experienced network administrators in more central topological positions, providing the basis over which to construct the inter-organizational relationships needed for the practice of network administration in general, and inter-domain routing in particular.

In addition to being channels between the global context of network operations, and the regional context of South Asia, instructors and those who make up the core of SANOG are themselves symbols of the global community of practice of network administrators. Instructors from outside South Asia are visibly foreign in appearance, manner and accent, representing an image of network administration in other regions of the world. Those who make up the core of SANOG have identifiable local origins, and represent the image of what it means to be a network administrator who is able to move between the South Asian context and more global contexts. When they talk about going to APRICOT or NANOG meetings, they communicate the image of meetings similar to SANOG in other parts of the world, and construct an image of themselves as being the kind of individuals who can meaningfully participate across these contexts.

I came to occupy a similar position when I mentioned in conversations that I had attended NANOG. This often changed the perspective of those I spoke with, as I became no longer just a student from the University of California, Berkeley, but also the kind of person who attends NANOG meetings. Overbreakfast at SANOG22, a SANOG committee member told me of his experience at APRICOT meetings, talking about how easy it was to approach senior network administrators and talk with them there. When I mentioned that I’d had a similar experience at NANOG meetings, he responded that he would very much like to go to a NANOG meeting, and hoped that he could get a scholarship to support his travel and board to attend a NANOG meeting. When I told him that I’d been to four NANOG meetings, he raised his eyebrows, and wanted to know more about what it was like to attend a NANOG meeting. This led into a discussion of my research, which he encouraged me to present at SANOG [F-SANOG22:136-140]. Through this interaction, and others, my respondent’s perspective of me changed, from viewing me as an unknown, somewhat exotic, outsider, to thinking of me as the kind of person who could usefully contribute to knowledge and education at SANOG.

Community at SANOG is not just constituted in relation to a global sense of commu-

nity, but is also constituted as a regional community, spanning South Asia. This regional sense of community is created through a recognition of similar concerns which arise from performing network operations from peripheral topological positions. For instance, the cost of bandwidth is a constant concern, as most autonomous systems in South Asia rely on their upstream network providers to connect them to the content being requested by users of their networks. This leads to the creation of IXPs that can reduce overall bandwidth requirements by offering common services such as access to CDN cache devices, and by providing direct interconnectivity between local autonomous systems. IXPs are in many ways representative of the mutual trust relationships which must be formed between autonomous systems to make them effective, reflecting a growth of national communities around IXPs. Equally, shared perspectives and concerns around the creation and operation of IXPs from peripheral positions provide the means for the establishment of relationships and community spanning countries in South Asia.

While there are strong trust relationships, constitutive of thick social relationships, amongst the core group of senior technical personnel at SANOG, the larger set of individuals who attend SANOG are much more diffusely connected to one another. Outside the core group, the form of community at SANOG is largely based on similar concerns and shared symbolic resources, rather than the closely knit relationships and mutual commitments of NANOG. This is not to say that shared symbolic resources and similar concerns do not play a role at NANOG, but rather that they become more visible at SANOG in the absence a dense network of social relationships. Accordingly, the form of community at SANOG is weaker compared to that at NANOG. The core group at SANOG is the anchor for this community. If it were to cease to be, then it is likely that SANOG would also cease, in the absence of individuals who can command the respect of technical communities in their countries and the rest of the world, and act as channels between global and local contexts. This form of community makes periodic interaction in shared physical space over a defined time all the more important, as it is only through these shared physically co-located moments that the symbolic content, concerns and relationships which construct a regional community can effectively be formed and maintained.

7.2.5 Extending Interconnection

All of the elements involved in extending the Internet that I've discussed so far - relationships, practice, community - ultimately operate in service to the goal of extending the interconnection of autonomous systems in the inter-domain routing system to South Asia. With the exception of a few large autonomous systems in India, all autonomous systems in South Asia rely on transit providers to connect them to the global Internet. The large Indian autonomous systems manage a proportion of these transit arrangements, especially for Bhutan and Nepal, which are landlocked. Interconnections between autonomous systems within a country are managed through a variety of arrangements, sometimes depending on private peering, sometimes relying on a larger autonomous system to act as a country-level transit

provider, and sometimes working through IXPs, which themselves operate under differing policies.

Nepal and Bangladesh both have IXPs developed by local ISP associations: the Nepal Internet Exchange (NPIX) established in Kathmandu in 2002, and the Bangladesh Internet Exchange (BDIX) established in Dhaka in 2004.²⁶ The development of both these IXPs was led by recognized senior technical personnel from these countries, who are part of the core of SANOG, and remain involved in the operation of the IXPs.²⁷ These IXPs were created with support from PCH and NSRC, which have been involved with technical training and IXP construction in many countries around the world. Apart from local interconnection, NPIX and BDIX also offer other services to their members, including hosting local DNS root server instances, and CDN cache devices. Network administrators from Nepal and Bangladesh spoke of their IXPs with pride, and talked about how the training they received at SANOG would be used to enhance the services these IXPs offered. For instance, a senior network administrator from Bangladesh told me how the knowledge of IPv6 that personnel from Bangladesh gained through the SANOG18 routing workshop would be put to use to expand the number of autonomous systems interconnected over IPv6 at BDIX. NPIX and BDIX represent outcomes driven almost entirely from within the Internet's distributed system of governance, drawing on local technical personnel and corporations, and non-profits operating in service to the Internet's global technical infrastructure. They are at once created by the extensions of relationships, practice and community into South Asia, and act as institutional and physical anchors for relationships, practice and community within their home countries.

Some ISPs in Nepal obtain their connectivity to the Internet via satellite. However, for high speed access, they must connect over terrestrial optical fiber cables, which are routed through India, since Nepal is landlocked. This is done by leasing capacity on a cable from an Indian telecommunications provider, or by paying an Indian autonomous system for transit to the global Internet. A consortium of ISPs in Nepal employs the former strategy to lease bandwidth across links that connect them to the Hong Kong Internet Exchange (HKIX), one of the most highly connected IXPs in Asia [F-SANOG18:182]. Bangladesh has its own coastline, and has direct access to submarine optical fiber cables, and therefore has no need to connect through India, although some ISPs in Bangladesh do maintain connectivity through India, if only to avoid routing through a third country to interconnect Indian and Bangladeshi autonomous systems.

Bhutan is too small a country to warrant an IXP. International connectivity for Bhutan flows through the state-run telecommunications provider, Bhutan Telecom, which in turn provides connectivity to several privately-run ISPs. Bhutan Telecom contracts with large autonomous systems in India to connect to the global Internet. Bhutan Telecom faces various challenges in maintaining international connectivity, not the least of which is the difficulty

²⁶Details of NPIX are available at <http://www.npix.net.np>; details of BDIX are available at <http://www.bdix.net>, last retrieved May 5, 2014.

²⁷The Nepali network administrator who was primarily responsible for the development of NPIX told me how he still maintains the BGP route filters at NPIX on a volunteer basis, updating them every few months to ensure that they are current[I26:3].

it faces in negotiating interconnection agreements, lacking the scale of many other country-level autonomous systems. For instance, a senior network administrator from Bhutan told me of the challenges they faced in obtaining CDN cache devices. In general, CDNs require that autonomous systems hosting their cache devices request a certain level of aggregate bandwidth from the CDN. Bhutan Telecom cannot satisfy this condition, driving up their bandwidth costs for interconnection through Indian autonomous systems. Thanks to conversations with Google personnel attending SANOG, Bhutan Telecom was able to obtain a Google CDN cache device, but were still unable to obtain cache devices from other CDNs, such as Akamai [F-SANOG22:39]. As the most peripheral country in the South Asian Internet - and indeed, on the global Internet - Bhutan's Internet presence is supported in no small part by personal relationships, such as those they established with Google personnel at SANOG meetings.

The IXPs in India and Pakistan are both operated with involvement from the state. I will discuss the Indian IXP, NIXI, in the next section. The Pakistan Internet Exchange (PIE) is the principal IXP in Pakistan, with locations in Islamabad, Lahore and Karachi. It is operated as a subsidiary of the Pakistan Telecommunications Company, Ltd (PTCL), which is largely owned by the state, but currently operated by Etisalat, a telecommunications company with global holdings headquartered in the United Arab Emirates.²⁸ PTCL also provides the bulk of Pakistan's international connectivity, managing landing stations in Pakistan for several submarine cables. Unlike Bangladesh, Pakistani autonomous systems do not have the option to connect through India, due to the long-running political ill will between Pakistan and India.

As these cases suggest, there are a wide variety of administrative arrangements for the management of interconnectivity across South Asia, affected by a range of factors: physical geography, market organization, the strength of the local technical community, the propensity for state intervention, and international relations.

Many of the IXPs in the South Asian region provide updates on their development at SANOG meetings. At NANOG meetings, IXP updates occur during the peering coordinators meeting, which is not recorded, and at which only a subset of NANOG attendees are present. In contrast, IXP updates at SANOG occur as part of the main conference program, and include NANOG-style "peering personals" for autonomous systems who wish to advertise their willingness to peer. At SANOG22, for instance, the IXP updates session included updates from BDIX, NIXI and NPIX, as well as "peering personals" from international and South Asian autonomous systems, including Limelight (a CDN provider), Google, Microsoft, PCH, Fiber@Home (a Bangladeshi ISP), HonestyNet Solutions (an Indian hosting provider), and Bhutan Telecom. The IXP updates provided brief histories of the IXPs, indications of future development - such as anticipated new members, services, facilities and IPv6 migration - and a statement of IXP locations,²⁹ size of membership, amount of traffic exchanged, and

²⁸See http://www.ptcl.com.pk/pd_content.php?pd_id=48 and <http://www.etisalat.com/en/about/profile/company-profile.jsp>, last retrieved May 5, 2014.

²⁹As IXPs grow, they may come to occupy multiple geographically separated facilities

additional facilities offered (such as CDN caches and hosted root server instances). The “peering personals” were much briefer presentations, including a statement of autonomous system number, peering policy and their points of presence.³⁰ Many presenters specifically asked the audience to contact them during the breaks in the conference to discuss peering. During his presentation, a SANOG committee member took a few extra minutes to show his autonomous system’s entry in PeeringDB, and actively suggested that all South Asian autonomous systems should have updated entries in PeeringDB [F-SANOG22:245]. Although the IXP updates session at SANOG is much sparser than that at NANOG, it does provide the basis for a regional sense of interconnectivity, and opportunities to propagate recognized best practices in interconnection (such as maintaining a PeeringDB entry), and offers possibilities for interconnection with autonomous systems with an international presence.

Personal relationships remain integral to the construction of the inter-domain routing system in South Asia, especially those held by the individuals who act as bridges between autonomous systems in South Asia, and the rest of the world. I was told repeatedly by those I spoke with how specific individuals, who occupy the “top level” at SANOG [I29:4] were able to help resolve interconnectivity issues. One of these individuals told me how he managed to establish IPv6 connectivity for the SANOG18 meeting in just a few hours, by leveraging his personal relationships. Similarly, he was able to help quickly clear filters blocking BGP routing announcements for a new address block allocation:

A few years ago we had a SANOG in Karachi. A bunch of Bangladeshi ISPs had received a new address allocation from APNIC. They come to Karachi, and they couldn’t access their routers from the SANOG network. We realized that FLAG³¹ was doing some filtering, because this is a new address block. The CyberNet guys who were providing transit for SANOG were trying to talk to FLAG and raise a ticket. Then I asked, what is your ticket number? So I got this friend on IM [instant messaging] and I said, hey, can you look at this ticket number and fix this problem? It got solved within minutes because she was the highest level engineer within FLAG in London, the two of them there, they fixed it. Then a few hours later, my guy got a call from FLAG saying, how did you escalate? [laughs] So those things still can happen. I think a lot of people realize that. That helps. [I26:7-8]

As this story illustrates, it can be difficult to establish the visibility of routing announcements across the Internet, since network reachability can only be evaluated from specific situated positions. If the Bangladeshi ISPs had not tested their network reachability from Karachi, it would likely have been a while longer before they realized that FLAG was filtering their newly allocated address space. This story also serves to remind us of the role of

³⁰The presenter for Fiber@Home, for instance, indicated that they operated an open peering policy, with points of presence in Mumbai, Chennai, London and Singapore.

³¹Fiber Link Around the Globe, a submarine optical fiber cable system providing connectivity to landing stations around the world.

RIRs: APNIC provides critical Internet resources to autonomous systems in the Asia-Pacific region, and maintains a presence at SANOG. APNIC staff are available for consultations on resource allocation, run workshops on resource management, and present an update during SANOG meetings on their activities and the usage of critical Internet resources in the Asia-Pacific region. At SANOG22, an APNIC representative talked about growth in the use of IPv6 in the region, policies for IPv4 allocations with the impending exhaustion of available IPv4 address space, inter-RIR transfers of resources and more [F-SANOG22:225-229]. The presence of APNIC at SANOG provides access to the system of institutional anchors which act as stewards of critical Internet resources, helping to maintain the stability in the topology of interconnections amongst autonomous systems making up the inter-domain routing system.

The stability of the inter-domain routing system does not just depend on the stability of individual interconnections amongst autonomous systems. It also relies upon autonomous systems to be well-behaved in the ways that they announce their address space, by not disaggregating without need, and being careful maintain stable routing announcements.³² Every SANOG meeting features an update on the state of the default-free routing table, with a focus on the behavior of autonomous systems in South Asia.³³ Autonomous systems which are behaving badly, by announcing maximally deaggregated address space, or by issuing large numbers of routing updates, are specifically named in order to try and change their behavior. The presentation is used to make the point that these individual bad behaviors have the effect of harming the entire Internet [F-SANOG18:191, F-SANOG22:205].

The process of disembedding relationships, practices and community from remote contexts, and then re-embedding them in contexts across South Asia, results in distinctive articulations of abstract global concepts. IXPs are formed and governed in very different ways across the different countries in South Asia, some with greater government involvement, and others formed through local ISP associations, with leadership from individuals who have relationships connecting them to global contexts. Relationships are formed in the process of learning practice, rather than in the process of engaging in practice. Community is constructed over shared concerns arising from peripheral topological positions and shared symbolic resources relating to the practice of network administration. Practice is taught with support from non-profits who engage in training activities for network administrators across the world. The extension of relationships, community and practice for inter-domain routing to South Asia is made possible only with substantial support from individuals and organizations with connections to more global contexts.

This support stems in part from deeply held beliefs of how network administrators around the world need to support one another in their everyday practice, especially through the vehicle of meetings such as SANOG and NANOG. This support also stems from a more pragmatic perspective: if no-one governs the Internet, then everyone is responsible for its

³²See chapter 4 for a more detailed discussion of these issues.

³³These presentations draw from reports updated daily at the CIDR Report, available at <http://www.cidr-report.org/as2.0/>, last retrieved May 7, 2014.

well-being. When network administrators at autonomous systems in South Asia engage in bad practices - such as disaggregation of their IP address prefixes, or hijacking others' address space - they have the potential to harm the Internet as a whole. In a distributed system of governance, those in more topologically central positions may wield more power of a certain kind by virtue of their position, but they are unable to exert absolute control over those at the periphery. The only path towards maintaining order in such a system is to ensure that individuals in any position - whether central or peripheral - abide by certain standards of practice, and maintain relationships and spaces through which to share information, coordinate and collaborate in service of the system as a whole.

The production of the Internet in South Asia depends upon global systems of practice, institutions and interconnections, yet manages to articulate these in distinctive forms, and still integrate with them to form part of the global Internet. In the following section, I discuss these articulations in the context of the Internet in India, focusing on the role of the state and market organization in the manner in which IXPs are constructed.

7.3 State, Market and Internet

The Internet's distributed system of governance does not operate in a vacuum, ruled by its internal logic. Thus far, I have focused largely on its internal dynamics, paying attention to market organization and state concerns only insofar as they impinge on this distributed system of governance. In this section, I switch this perspective, to more explicitly examine the role of the state and the market in the production of the Internet, especially in relation to one of the principal means of interconnection of autonomous systems, the IXP.

In most parts of the world, IXPs are community-run non-profit entities, or privately provisioned. Although an IXP typically operates in only one physical location, some larger IXPs may operate across multiple locations in the same city, or even in different cities altogether, sometimes in competition with other IXPs servicing the same region.³⁴ In contrast, there is only one IXP provider in India, the National Internet Exchange of India (NIXI), constituted as a public-private partnership, operating in several cities across India. In comparison to most IXPs, this is an unusual model of operation.

As a distinctive local articulation of the concept of the IXP, NIXI offers a critical site at which to examine the interplay of state and market with the Internet's distributed system of governance. Why was it that ISPs in India were unable to create IXPs of their own? Why was the involvement of the Indian state needed to form an IXP in India?

I argue that the operation of NIXI as a public-private partnership is the result of the combination of an interventionist state, and an oligopolistic market structure. Although the

³⁴For example, Equinix, a private provider of IXP (and other) services, operates IXPs in multiple locations across 19 cities in Europe, North America and the Asia-Pacific, see <https://ix.equinix.com/ixp/ixLocations>. In Singapore, Equinix competes with two other IXPs, which both operate as non-profit entities: the Singapore Open Exchange (see <http://www.sox.net.sg/>) and the Singapore Internet Exchange (see <http://sgix.sg/en/>). All URLs last retrieved May 8, 2014. See chapter 6 for a discussion of IXPs.

Indian economy was liberalized in the 1990s, the Indian state has long operated on a philosophy of self-reliance and strong regulation, building a state-run industrial base, with a highly regulated private industry. The market for ISPs in India is split between four large ISPs with national, and international, infrastructure, and many much smaller ISPs. I will show how the combination of these factors created the conditions under which state involvement was necessary for the creation of IXPs in India. This was by no means a straightforward process; it was messy, proceeding by twists and turns, and involving unanticipated consequences.

7.3.1 The Origins of the Internet in India

As in the USA, the state played a strong role in the development of the Internet in India. However, this role was not always supportive. Some parts of the Indian government acted to further the development of Indian Internet infrastructure, while others were threatened by these developments, and acted to curb them. Like any state, the Indian state is not a monolithic entity with singular intentions, but is rather composed of multiple competing components, often with differing interests.

Unsurprisingly, the emergence of the Internet in India was closely tied with the development of the Internet in the USA. Internet technologies arrived in India through a combination of the circulation of people between India and the USA, and the acquisition of computing systems which used TCP/IP as their default mechanism for networking. Indian research institutions were working on computer networking technologies as early as the 1970s. Indians educated in the USA in that period returned to India with knowledge of the packet-switching technologies developed for the ARPANET, and proposed the use of similar technologies in the development of Indian computer networks. The first Indian packet switching networks were powered by domestically produced hardware and software. As packet switching matured in the USA, computers from American corporations, such as DEC, became available for use in Indian computer networks in the 1980s. These computers ran on a Unix variant derived from BSD which had a TCP/IP networking stack and higher level network services built in.³⁵ A nationwide computer network connecting research institutions, the Education and Research Network (ERNET) was developed in the 1980s using TCP/IP networking technologies. This effort was led by computer science researchers in India, supported by bureaucrats at the Indian government's Department of Electronics.(Ramani 2011)

Even with support from the Department of Electronics, it was sometimes difficult to work around the highly regulatory nature of the Indian state. For instance, the use of TCP/IP protocols was forbidden over leased lines³⁶ in the early 1980s, since the Indian government required the use of international telecommunications standards set by the CCITT.³⁷ Links

³⁵The Berkeley Software Distribution (BSD) was one of the earliest packaged distributions of Unix. As the name suggests, it was maintained by a computer science research group at the University of California, Berkeley.

³⁶Leased lines are dedicated telecommunications links leased from a third party provider.

³⁷The Comité Consultatif International Téléphonique et Télégraphique (CCITT) was an international treaty organization which set standards for telecommunications. It later became the ITU-T, the telecom-

between ERNET nodes were initially operated over dial-up links, shifting to leased lines once the the regulations curtailing the use of TCP/IP were repealed. The Indian government was able to put such regulations in place in part because domestic telecommunications infrastructure was originally constructed and operated by the Indian government's Department of Telecommunications, which at the time functioned as both regulator and operator for telecommunications in India. As a computer scientist involved with the project at the time put it, "We worked quietly, like a resistance movement, to survive restrictive telecom policies, and a lack of good telecom infrastructure in many places." (Ramani 2011:57-58)

Responsibility for domestic telecommunications infrastructure was spun out to a government-run corporation, Bharat Sanchar Nigam Ltd (BSNL) in 2000 (Chowdary 2011:68).³⁸ Prior to this, another government-run corporation, the Mahanagar Telephone Nigam Ltd (MTNL), was created in 1986 to provide improved telecommunications services in major metropolitan areas: for New Delhi, as the capital of India, and for Mumbai, as India's financial center.³⁹ When BSNL was created, it was given responsibility for the rest of the country.

Computers on the ERNET were first connected to the international locations for email services in 1986, through a link to a server located at a research institute in Amsterdam, which was in turn provided connectivity to computers across the world. This link was provided by Videsh Sanchar Nigam Ltd (VSNL), a corporation created by the Indian government for the purposes of managing connectivity for international telecommunications.

This was the environment for connectivity which the nascent Indian software industry had to contend with in the early 1990s. Since BSNL had yet to be formed, the software industry had to obtain domestic telecommunications links from the Department of Telecommunications in order to connect to the international gateways operated by VSNL. To get around the difficulties of negotiating domestic links with the Department of Telecommunications, those involved with the software industry allied themselves with supporters in the Indian government's Department of Electronics, to create a new vehicle, the Software Technology Parks of India (STPI). STPI is better known today for the role it played in obtaining tax breaks to support the emergence of the Indian software industry. However, an industry leader who was involved with this process told me how one of the primary original intentions of STPI was to act as a legal vehicle through which domestic connectivity to VSNL's international gateways could be independently obtained, bypassing the regulatory limitations and bureaucratic overhead of the Department of Telecommunications [I31:2-3].

All domestic telecommunications infrastructure and international links were managed by entities connected to the Indian state in this period. The development of the early Indian Internet had to proceed over infrastructure provided by these entities, and so was subject to control and regulation by these entities. Both government-supported research efforts, and private industry, attempting to obtain connectivity to the Internet had to work through a

munications standardization sector of the International Telecommunications Union.

³⁸Until the liberalization of India's economy in the 1990s, the majority of infrastructure in India - including telecommunications - was provided directly by government ministries, or by "public sector units" (PSUs), which are government-run corporations operating, for the most part, as monopolies in their areas of concern.

³⁹From <http://www.mtnlmumbai.in/index.php/about-us>, last retrieved Jun 10, 2014.

highly regulated monopolistic environment. They did so by working with supporters within the Indian government to create institutional contexts through which alternative regulatory policies could be established. These new institutional contexts provided the means through which the topological arrangements required for Internet connectivity could be established, whether over existing Department of Telecommunications links (as in the case of ERNET), or through independent physical connectivity (as in the case of STPI).

In 1992, following the Indian government's moves towards economic liberalization, the Department of Telecommunications introduced a policy for "Value Added Services", which operated over domestic telecommunications infrastructure. Licensed private operators were not allowed to provide generic Internet access, but were rather limited to the provision of specific services, including email and FTP. Matters were made more difficult by the fact that VSNL competed by offering similar services, while operators were required to connect through VSNL's international gateway. In response, operators of these services banded together to form the Email and Internet Service Provider Association of India (EISPAI), to lobby the Department of Telecommunications for changes to these policies. In 1997, an independent telecommunications regulator was created, the Telecommunications Regulatory Authority of India (TRAI), which took over the Department of Telecommunications' regulatory authority. TRAI and the Department of Telecommunications soon came into conflict over who was to have control of Internet policy in India. Eventually, it was through the actions of the Prime Minister's Task Force on Information Technology and Software Development that a policy for licensing privately operated ISPs was put into place, in 1998. Shortly thereafter, EISPAI renamed itself ISPAI: the Internet Service Providers Association of India (Singhal 2011).

The Indian government was critical to the development of the Internet in India, as a regulator and holder of a telecommunications monopoly, and as an enabler of those who would challenge this situation to provide Internet services in India. Then, as now, the Indian government was the key site which had to be negotiated in order to develop the Indian Internet. In the process, the monopoly power of the Department of Telecommunications was gradually broken down into an independent regulator (TRAI), government-run corporations (MTNL, BSNL and VSNL), and privately-run ISPs.⁴⁰ In the following sections, I discuss the nature of the private and public arrangements for telecommunications in more detail, through an examination of the organizational forms required for interconnection of ISPs in India.

7.3.2 The Creation of NIXI

Even though the market for private ISP operations in India was opened up in 1998, the first IXP in India was only constructed in 2003, through the creation of the National Internet Exchange of India (NIXI). The genesis of NIXI was the result of parallel development of the

⁴⁰Several ISPs have since become broad-based telecommunications providers, offering Internet access alongside cable television, and wired and cellular phones.

idea of an IXP for India by the Department of Information Technology and ISPAI. The interest of the Indian government in an IXP was to keep domestic traffic within Indian borders, partly as a means to support the domestic ISP industry, and partly from security concerns over the possibilities for the surveillance of Indian data traffic carried outside India [I32-1:10]. ISPAI considered the development of an IXP a natural extension of their position as an ISP association, viewing the IXP as a service to support the activities of ISPAI members.

ISPAI faced two principal challenges in their attempt to form an IXP. First, they lacked the technical capacity to develop an IXP: they were a small organization with just three or four full-time staff, created for the purposes of lobbying for the nascent Indian ISP industry. Second, they were unable to raise the capital necessary for the creation of an IXP. When I asked a former ISPAI official why this was so, he responded that it was because of a lack of trust amongst Indian ISPs [I35-2:9]. This lack of trust was, and continue to be, due to the oligopolistic structure of the Indian ISP market, and the competitive focus of Indian ISPs. An official at NIXI confirmed this perspective, telling me how competitive concerns created an environment in which ISPs could not cooperate, requiring the government to step in as a neutral actor:

What I feel in India is that industries feel that if they put up a utility point, how will they come together? Even if they come together, where to set up the exchange point? That was the major hurdle. From their point of view, suppose they put it in one of the private ISP's place, then he will be knowing how traffic is going... So that's why they thought the government is neutral, because the government is not an ISP, so it will run better. [I32-2:4]

While the development of ISPs in the USA was led largely by technical personnel, the development of the Indian ISP market occurred through investments from old Indian business houses.⁴¹ In consequence, a strong Indian community of network operators - and the concomitant inter-organizational trust relationships - never developed in the manner that it did in the USA, or even in neighboring countries such as Nepal and Bangladesh.

As I noted earlier, Indian ISPs do not participate extensively in SANOG, which is the only organization promoting a network operator community within South Asia. Larger Indian ISPs choose to participate in international settings, which is where their concerns lie, since they need to establish relationships for operational and business purposes to maintain international connectivity. Small ISPs, operating in regional markets, typically lack the capital to participate regularly in SANOG. In addition, since they rely on large ISPs to carry their traffic domestically and internationally, they have little need to have their personnel

⁴¹According to the latest TRAI performance indicators report for the last quarter of 2013 (Telecom Regulatory Authority of India 2014:27-28), the largest overall provider of Internet service (combining wired and wireless markets) is Bharti Airtel with 23.37% of the market, followed by Vodafone with 19.13%, BSNL with 17.21%, Reliance Communications with 15.24% and Idea Wireless with 10.69%. Considering wired connections alone, BSNL is by far the market leader, with a 71.91% market share. I could find no similar data on Internet interconnection markets in India, although the size of the consumer market share is indicative of the reach of an ISP's national infrastructure.

engage in the system of interpersonal trust relationships which characterizes the Internet's inter-domain routing system. They may, however, send their personnel to SANOG meetings - when they are held in India - for training purposes. The Indian ISP market is characterized by large ISPs with international interests, and small ISPs with regional interests who are dependent on the large ISPs.⁴² This oligopolistic market structure - and associated quasi-hierarchical topological arrangement - alongside a culture of competition and government regulation, created a low trust environment in which it was difficult to cooperatively form an IXP.

Faced with an inability to raise capital from its membership, ISPAI approached the Department of Information Technology for the resources with which to form an IXP. This led to the creation of NIXI as a public-private partnership, albeit one in which the Indian government holds considerable authority. The Secretary for the Department of Electronics and Information Technology (DEITY) is the chair of the NIXI board.⁴³ The CEO of NIXI is an employee of DEITY, maintaining his primary office in the DEITY administrative complex. Of the 13 NIXI board members, four are in the employ of the Indian government its agencies, seven represent ISPs, one is the president of ISPAI, and one is an academic.⁴⁴ NIXI would not have come into existence without the involvement of the Indian government, which the Indian government parlayed into significant policy and operational control over NIXI.

7.3.3 Interconnection and Topological Organization

The Indian government's role in the creation of NIXI was no magic bullet for increased domestic interconnection. The same factors which created a low trust environment amongst ISPs in India also militated against a neutral interconnection environment. The most significant challenge NIXI faced in its early years was the reluctance of large ISPs to interconnect at NIXI locations. This was particularly an issue for VSNL, which controlled international telecommunications over the period of NIXI's creation.

Prior to the formation of NIXI, VSNL acted as the de facto IXP for Indian ISPs, as the principal entity which could carry traffic internationally, through which all ISPs had to be connected, even if indirectly through another domestic ISP. This situation changed as private ISPs began to establish international connectivity of their own. However, VSNL retained its topological position - and associated market power - as an international gateway for a few years more, as private ISPs took time to develop their own infrastructure and international relationships. The Tata Group, one of the major business houses in India,

⁴²The exception to this rule is BSNL, which is a large ISP with only domestic infrastructure. It is a special case, since it continues to be controlled by the Indian government.

⁴³The Department of Information Technology was renamed DEITY in 2012. The position of Secretary is amongst the highest that an Indian bureaucrat can hold, overseeing the activities of an entire department or ministry.

⁴⁴The list of NIXI board members is available at <http://nixi.in/en/about-us?id=114>, last retrieved Jun 17, 2014. Of the ISP representatives, three are from large ISPs (Tata Communications, Bharti Airtel and BSNL) and four are from medium-sized ISPs operating in multiple regional markets.

eventually acquired a 45% stake in VSNL in 2002 as a means towards creating their own ISP business. They went on to acquire several other global network providers to ascend to Tier 1 status as Tata Communications.⁴⁵

VSNL initially refused to interconnect at NIXI, perceiving only a threat to its dominant topological position. It was eventually persuaded to interconnect at NIXI once a formula for the settlement of interconnection fees favoring VSNL was put into place. Unlike many other IXPs, NIXI manages the settlement of fees amongst connected ISPs; in most other IXPs, settlement fees are typically a matter of bilateral agreements amongst ISPs, or set through participation in a multilateral peering agreement. Settlement fees at NIXI are calculated using the so-called “x-y” formula.⁴⁶ This is based on a “requester pays” model: the volume of data requested by one ISP from another is subtracted from the volume of data requested in the opposite direction, and then multiplied by a fixed rate to arrive at the fee to be paid.⁴⁷ NIXI is the agency which calculates, collects, and settles these fees amongst connected ISPs.

Comparatively little content is hosted in India, compared to the rest of the world, largely due to problems with the infrastructure - particularly electrical power supply - required to operate data centers reliably. As a result, the majority of data flowing to consumer ISPs originates from outside India. The large autonomous systems which control international gateways are topologically positioned to reap the benefits of this structure, as consumer ISPs must depend on them to connect to the global Internet. The “x-y” formula essentially institutionalized these arrangements, protecting the revenue streams of autonomous systems with international gateways, such as VSNL.

The “x-y” formula is not the only mechanism through which the positions of large ISPs are institutionalized at NIXI. During an interview, a NIXI official, asked me how NIXI could be improved. I suggested that perhaps CDN cache devices could be installed at NIXI locations, to save bandwidth costs for small ISPs. The official responded that CDN cache devices were prohibited at NIXI, since these would erode the revenue of large ISPs, which continue to hold significant power over NIXI policy [I51:9]. Similar concerns prevent NIXI from providing direct interconnection between its various locations, since autonomous systems providing national backbones would then perceive NIXI as a competitor in long distance domestic connectivity.⁴⁸

In comparison to many other Internet governance organizations, NIXI’s internal operations are relatively opaque. When I asked a NIXI official if I could have access to meeting minutes of the NIXI board, he responded that he would not be allowed to provide me with those materials. He told me how he had tried on several occasions to open up NIXI’s policy development process, but that he was confounded by a combination of the power held by

⁴⁵See <http://microsites.tatacommunications.com/about/history.asp>, last retrieved Jun 13, 2014.

⁴⁶This was the term favored by all those I spoke with about NIXI.

⁴⁷See NIXI’s Routing and Tariff Policy, available at <http://nixi.in/en/routing-and-tarrif-policy>, last retrieved Jun 15, 2014.

⁴⁸NIXI currently has 7 locations in major metropolitan areas in India, at Mumbai, New Delhi, Chennai, Kolkata, Bangalore, Hyderabad and Ahmedabad. See <http://nixi.in/en/noc-locations>, last retrieved Jun 16, 2014.

large ISPs on the NIXI board, and apathy from small ISPs who could not perceive any benefits from their involvement [I51:11]. NIXI officials are constantly concerned with the amount of traffic which transits their facilities, since this is much lower than they expect. In past years, they contemplated having the Indian government mandate interconnection at NIXI for all ISPs, but abandoned this approach as they realized that this would be prohibitively expensive for many small ISPs, who would have to lease capacity on dedicated network links to carry their traffic to a NIXI location [I35-2:23].

IXPs are often intended to disintermediate the power of larger autonomous systems, by providing direct interconnection between smaller autonomous systems. Paradoxically, NIXI instead mirrored and reinforced the market power of large autonomous systems. It would have been hard for NIXI to do otherwise if it were to succeed in persuading large Indian ISPs with international connectivity to interconnect through it. The dilemma that NIXI faced was the result of policies which favored monopolies over specific areas of domestic and international connectivity, and a global environment in which the majority of content was hosted outside India.

7.3.4 Internet Governance in India

As NIXI matured, it developed into the principal agency for Internet governance activity in India, incorporating responsibilities for the management of names and numbers into its functions. ISPAI did consider undertaking similar activities, but bowed out in favor of NIXI. Through these processes, NIXI shifted from largely domestic concerns over interconnection, to being formally tied into global Internet governance institutions. It was able to do so by claiming its position as the legitimate representative of the Indian government, while at the same time remaining tied to private industry within India.

NIXI took over the “.in” ccTLD registry in 2005, from the Center for the Development of Advanced Computing (C-DAC), which had previously held this responsibility as the nodal agency for the development of ERNET.⁴⁹ Accordingly, NIXI formed institutional ties with ICANN for this purpose. NIXI created the Indian Registry for Internet Names and Numbers (IRINN) in 2012 to allocate IP address space and autonomous system numbers within India. It did so through a National Internet Registry (NIR) agreement with APNIC, under which NIXI is allowed to set policy for the allocation of critical Internet resources within India, but cannot form any policies which contradict APNIC policies. With respect to the management of names and numbers, NIXI is beholden to institutions which, to varying degrees, serve the Internet’s technical communities. NIXI personnel regularly attend meetings of various Internet governance institutions, as do NIXI fellows, who are Indian citizens funded by NIXI to increase Indian participation in international Internet governance activity.⁵⁰

⁴⁹See <https://registry.in/Policies/C-DAC%20Closure>, last retrieved Jun 16, 2014.

⁵⁰At SANOG18, a couple of senior technical personnel commented on the size of the Indian delegation to the last APNIC meeting they had attended, totaling about 40 people, at least half of whom were sponsored by NIXI [F-SANOG18:185].

This curious inversion of power - in which private authority formally supersedes state authority - is at the center of many recurrent debates around Internet governance. NIXI presents the face of the Indian state to this system, even though it is, at least in part, controlled by the interests of large Indian ISPs. I first met NIXI officials when they were in the process of formulating IRINN in 2011, and grappling with their understandings of the technical and institutional issues involved. A NIXI official told me how they decided to form IRINN for two principal reasons. First, to support the domestic ISP industry, which at the time had to spend foreign exchange at APNIC in order to obtain number resources. Second, to support the activities of the Indian government, which was developing infrastructures such as the National Knowledge Network⁵¹ and the Aadhar card.⁵² A NIXI official I spoke to went on to suggest that this was also a point of national pride; as a regional superpower, India should have a NIR, since several other countries in the Asia-Pacific region - China, Japan, South Korea and Vietnam - all have NIRs of their own.

NIXI's formal ties to APNIC were negotiated through the process of forming IRINN. It rankled NIXI officials that APNIC would not give them a contiguous block of IPv4 address space for their needs, but would rather only provide multiple smaller fragmented IPv4 address blocks [I32-1:8-9].⁵³ When I interviewed NIXI officials after IRINN had been formed, they spoke with pride of the technical demonstration of the IRINN site to APNIC officials, their subsequent certification as the Indian NIR by APNIC, and the success they had achieved in having hundreds of domestic ISPs sign up with IRINN for number resources [I32-2:6].

Through the process of taking on the “.in” registry, and creating IRINN, NIXI developed additional internal capacities, and external institutional links, to evolve beyond its initial role as an IXP provider. It came to perform multiple Internet governance roles in interconnection and resource allocation. In doing so, it acted as an institutional anchor to provide services in a low-trust domestic environment, and became the institutional face of India to the international institutions of Internet governance.

The Internet's distributed system of governance relies on strong technical communities being able to exert their power “for the good of the Internet” in relation to nation states and corporations. This governance system reaches its limits in the absence of a strong technical community, at which point mechanisms provided by nation states and market organization are required to fill this gap. Yet it may be that the very causes of these mechanisms - which functionally substitute for a strong technical community - inhibit the development of a technical community in the first place. The possibilities for the formation and continuation of strong technical communities are bounded by their economic and political environment, and by the topological positions of the autonomous systems represented by a community in

⁵¹A high speed network intended to connect thousands of educational and research institutions across India.

⁵²A program to create a national identification card for Indian citizens. My interviewee suggested that each card could potentially be linked with a unique IPv6 address.

⁵³This was at a time when IPv4 address space was already running out, and APNIC was allocating new IPv4 address space under special policies to manage scarcity. See <http://www.apnic.net/community/ipv4-exhaustion/ipv4-exhaustion-details>, last retrieved Jun 18, 2014.

the Internet's inter-domain routing system.

The norms which guide market and state governance mechanisms - profit, or control - may be quite different from those of trust, coordination and collaboration, which guide the Internet's technical communities to provide a balance against economic and political interests. The resulting local articulations of Internet governance are distinctive, yet still integrate with the global system of Internet governance through institutional and interpersonal relationships which follow established structures, norms and practices.

7.4 Re-embedding the Internet

From a technological perspective, the Internet is a single global system based on common standards. However, the distributed nature of the Internet creates conditions in which governance arrangements may vary quite radically from one context to another. Everyday practices and governance structures are disembedded from their points of origin to imagine a global system of governance for standards, resources and topology. In order to have effects in the world, these elements of distributed governance must then be re-embedded in remote contexts. As they are re-embedded, they are also integrated into the Internet's global system of distributed governance; just as the Internet is a global system of interconnected networks, its governance system is similarly a global system of interconnected contexts. The process of re-embedding and integration may differ quite radically from one context to another, as I have shown in this chapter, through variations in the topological position of autonomous systems, and distinct historical and political-economic conditions. The Internet's distributed system of governance is composed of a unity of differences; a unity which is made possible only through significant ongoing social effort, which I characterize as spatial practice (Lefebvre 1991), the practice involved in the production of the space of Internet infrastructure.

In earlier chapters, I showed how trust relationships are integral to the maintenance of order in the Internet's inter-domain routing system, especially amongst topologically central Tier 1 and Tier 2 autonomous systems. I discussed the ways in which these relationships are maintained through the technical community which gathers at NANOG. SANOG is imagined as a South Asian analogue to NANOG, yet it could not be more different. Those who gather at SANOG represent the periphery of the Internet rather than the center. They form community amongst themselves through shared peripheral concerns, rather than to mitigate shared risks and uncertainties in inter-domain routing. The nature of community at SANOG is based on common symbolic resources, rather than the dense mesh of trust relationships which characterizes community at NANOG. Attendees at SANOG are in large part younger and less experienced than their counterparts at NANOG.

In spite of these differences, SANOG forms part of a system of network operators groups around the world, which aim to integrate network administrators into common understandings of practice, and provide the spaces for network administrators to make sense of the contexts within which they operate. These groups are linked by senior network administrators who provide local and global leadership: by organizing their regional network operators

groups, traveling to other network operators groups, and providing training. This system of network operators groups provides a linked set of localized anchors for the relationships, and shared practices and norms which structure the Internet's distributed system of governance. SANOG is different from NANOG because it serves a set of concerns which are those of the periphery of the Internet's inter-domain routing system (and of Internet infrastructure in general), rather than the center. However, SANOG and NANOG must not be understood only in terms of difference, but instead in terms of the ways in which they relationally constitute different positions within a single system of governance. They each do different work required of their specific position and context, while still connecting to one another to progressively disembed, re-embed and integrate practices and norms.

The space of the Internet's logical layer - created through the Internet Protocol - is formed and stabilized through the work of network operators groups, alongside the work of the centralized institutions of Internet governance. SANOG and NANOG represent points of variance which are linked together to produce this space. Yet the production of this space does not occur in a vacuum. At every moment of production, of re-embedding, it rubs up against the territorial space of the nation state, and the economic motivations of corporations, both of which may militate against the coordination and collaboration which characterize the Internet's distributed system of governance. Under certain conditions, such as those in India, the formation of local technical community may be retarded. In such situations, alternate governance arrangements are required to provide the functional outcomes - such as shared infrastructure for IXPs, or the local provision of critical Internet resources - which might otherwise be provided by a strong technical community.

The case of NIXI, in which the Indian government was involved in the formation of an IXP, is instructive to both understandings of how the territorial space of the state is produced, and of how the state may become involved in the production of the space of the Internet. By allowing the licensing of private ISPs and privatizing state-run telecommunications corporations (as with VSNL), the state partially de-territorialized, by giving up absolute control of telecommunications infrastructure. By providing the institutional structure and technical infrastructure for NIXI, the Indian government re-territorialized, acquiring a degree of control over a critical element of Internet infrastructure in India. This process involved realignments within the structure of the government, as the Department of Telecommunications was gradually relieved of monopoly over telecommunications infrastructure, and the Department of Electronics and Information Technology took on responsibility for Internet governance activity through the vehicle of NIXI. Through NIXI, the Indian state integrated itself into the Internet's distributed system of governance: providing support for SANOG meetings, managing Indian representation at APNIC, and tying into the centralized institutions of Internet governance for the domestic provision of critical Internet resources.

The territorial space of the state must, therefore, be understood as a space which is actively produced and negotiated through ongoing interventions, rather than as a passive physical space bounded by borders. Territorial space is perhaps better comprehended as a networked form in which the networks of state authority are formed through interlinked physical infrastructure and institutional structures. Once the territorial space of the state

is conceived as a network, a clearer perspective can be formed on how the space of the state integrates with the space of other infrastructures, including the Internet. Through this process of integration, the state may become a vehicle for the interests of the Internet's technical communities - for instance, by forming a NIR in conformance with APNIC policy - just as much as the state forms new capacities to extend itself into broader supra-national contexts of Internet governance. By reflecting on the interaction between the state and the Internet's distributed system of governance, it is possible to gain a clearer perspective on each of these systems of power, as the state is produced in the Internet, and the Internet is produced in the state.

The Internet's distributed system of governance allows for substantial internal variation, integrating a variety of interests across different scales of operation. This is a remarkable feat, which speaks volumes of the Internet's openness and capacity for change, while still providing a unified global infrastructure. As the Internet's distributed system of governance is renegotiated in its encounters with state and market interests, it faces increasing challenges to its norms of openness, trust, coordination and collaboration. The outcomes of these negotiations may vary substantially across different contexts. This raises a critical question, which I will attempt to address in the following chapter: to what degree can these negotiations disembed themselves from their contexts to alter the fabric of the Internet's global system of governance?

Chapter 8

Conclusion: For the Good of the Internet

When the Internet is spoken of, it is often framed as something wholly new - a break with the past, offering new possibilities for society - or as a continuation of the means of control of established state and market powers. In this dissertation, I uncover that which is novel in the Internet, while placing this novelty in relation to that which is a continuity with the past. To do so, I examined the “logical layer” of the Internet, which mediates between the novel appearance of democratized virtual space that the Internet provides, and the physical space of telecommunications infrastructure, the realm of traditional powers. Rather than accepting that the Internet provides new political opportunities for “freedom” and “democracy”, I sought to ask why and how these opportunities are made possible in the first place. It is critical that we ask this question if we are to retain the opportunities that the Internet provides: we must understand what features of the Internet provide the possibilities for these opportunities, so that we may learn from, protect, and enhance these features.

I addressed this question by examining the Border Gateway Protocol (BGP), the technology which is used to manage the interconnections amongst networks - the inter-domain routing system - which make up the global Internet. This examination led me in many directions, from studying the origins of BGP in the early history of the Internet, to understanding the risks and uncertainties inherent in BGP, to uncovering the social formations and practices involved in the operation of BGP. The results of my research indicate a novel form of governance involved in the operation of BGP, and indeed, of Internet infrastructure in general, which I have called “distributed governance”. If the Internet is a distributed system which can “route around failure” (and around established forms of power), as the saying goes, then the arrangements involved in the governance of this system must themselves, surely, be distributed, to avoid providing centers of power vulnerable to capture by established powers.

In the remainder of this chapter, I return to, and refine the three related propositions I outlined in Chapter 1, to articulate the nature of distributed governance.

First, that *trust* is the foundation upon which distributed governance operates: interpersonal trust, trust in institutions, assumptions of trust built into technological form, and arrangements of trust relations following the topological form of network interconnectivity. I address this first proposition in section 8.1.

Second, that trust is made possible only through physically co-located interactions *in space*. Extensions of trust relations *across space* require the formation of a network of similar spaces for physically co-located interaction, trusted individuals who can form the links between these spaces, and institutions to anchor these spaces. I capture the logic of these first two propositions through what I term the “practice of infrastructure”, which expresses the ways in which the everyday practice of managing network interconnection is premised upon the trust relations and spaces of the distributed governance of Internet infrastructure. I address this second proposition in section 8.2.

Third, that the engagement of established powers with the Internet’s distributed system of governance occurs through manipulations of space and trust, but is limited primarily by the technological form of the Internet which assumes trust. I argue that the power that the Internet’s technical communities hold is based upon the durable technological forms of infrastructure which produce the space of the Internet’s logical layer. If these technological forms were to change, then the possibilities for governance of these technologies would similarly change. This third proposition is a cross-cutting concern, addressed throughout this chapter, but primarily through the concept of “architectural embeddedness” in section 8.1.3, through the examination of the relationship between trust and security in section 8.1.4, through the analysis of the assumptions of power relations built into BGP in section 8.2.1, and through the examination of the interaction between state, market and technical community in 8.2.3.

Bringing together notions of trust and space, these three propositions illuminate what I call the “architecture of trust” that enables distributed governance.¹ It is by understanding the architecture of trust at play in Internet infrastructure that we can make sense of the nature of power in distributed governance: encoded in technological form, enacted in practice, embodied in technical communities and institutions.

The fundamental difference between established powers and distributed governance lies in spaces which are produced through conflicting logics. Each must change to accommodate the other, and it is through the examination of these processes of change in spatial production and trust relations that we may understand the nature of the Internet, and by association, the nature of the “information society”.

¹“Architecture of trust” and “trust architecture” are terms commonly used in computer science when talking about so-called “trustworthy computing” initiatives to create more secure computing systems. My objective in using this term here is to extend the concept of trustworthy computing to more thoroughly engage with trust as a social problem which requires social solutions alongside technological solutions; and which takes seriously the ways in which limits and possibilities for power are built into technological form.

8.1 Trust in Technology

The sense of “trust in technology” that I employ is threefold. First, as the role that trust plays in the governance of technological systems. Second, as the ways in which technology is perceived as trustworthy, whether in the abstract understanding of the Internet, or in the more specific understanding of inter-domain routing.² Finally, as the ways in which assumptions of trust are embedded into technological form. In my analysis, I combine these three senses of “trust in technology”, to understand how trust is the primary mechanism through which distributed governance operates, to mitigate the risks and uncertainties in inter-domain routing.

Interpersonal trust relationships are but one possible response to problems of risk and uncertainty, which may also be addressed through policing by centralized institutions (Gellner 1988), or through assurance structures (Yamagishi and Yamagishi 1994). Some argue that risk and uncertainty may be mitigated through market mechanisms. Following Polanyi (2001), I believe that markets are embedded in society, made possible only through the action of centralized institutions, such as regulatory agencies, or through trust relationships in market exchange (Granovetter 1985). Accordingly, I will limit my analysis of mechanisms for reducing risk and uncertainty to interpersonal trust relationships and centralized institutions. I will use these mechanisms to show how they act to maintain the embeddedness of markets for interconnection in inter-domain routing.

Interpersonal trust relationships and centralized institutions (or assurance structures) represent two extremes on a spectrum of possible arrangements for the reduction of risk and uncertainty to enable cooperation (Cheshire 2011). Trust is always interpersonal, a dyadic social relation. Accordingly, trust is never transitive: if A trusts B, and B trusts C, it does not follow that A trusts C (Christianson and Harbison 1997). Following the “encapsulated interest” account of trust from Hardin (2002), while B encapsulates A’s interests, and C encapsulates B’s interests, it cannot be said that C completely encapsulates A’s interests. However, risks and uncertainties can be transitive: A’s trust in B may in part be premised upon risks and uncertainties that B takes on through trust in C. These issues of transitivity are especially important to the analysis of large, complex systems, as risks and uncertainties which flow transitively across a system via technological form - as features of a system - cannot be addressed by trust relationships alone.

It is infeasible for everyone to form trust relationships with everyone else in large, complex social systems. In such systems, other mechanisms are required to enable cooperation. These may be mechanisms to evaluate the trustworthiness of unknown others, such as reputation (to evaluate individuals based on past behavior) or generalized trust (to assume a basic common level of trustworthiness). Alternatively, these may be mechanisms to warrant interactions, providing assurance structures that take on the risk and uncertainty in an interaction, allowing individuals to cooperate as though risk and uncertainty were absent.

²I use the general perspective of trust in the Internet by society at large to comment on trustworthy computing initiatives. However, the principal focus of my work remains upon trust in the centralized institutional structures and practices of the technical communities who operate Internet infrastructure.

In large, complex systems, such as the Internet's inter-domain routing system, networks of interpersonal trust relationships, and centralized institutions, work in tandem to provide an ordered, stable infrastructure. The balance between these arrangements - centralized vs. distributed - shapes the nature of governance of a system. These arrangements flow in part from historically articulated forms, and in part from the architecture of the technologies which these arrangements are intended to stabilize. In the case of inter-domain routing - and, I suspect, of most large scale modern infrastructures - governance arrangements and technological form evolved alongside one another, and shape each other.

Inter-domain routing is made possible through three related functional areas, each of which have different risks and uncertainties associated with them: topological arrangements for inter-domain routing, the management of critical Internet resources (such as IP address blocks and autonomous system numbers), and the development of Internet protocol standards. These represent the schema which I will use in the analysis which follows: *standards*, which specify the *resources* and interfaces, required for the creation of interconnections to form *topology*.

8.1.1 Trust in Design

BGP was created in the context of the NSFNET, a research network in which uncertainties were high as a result of ongoing changes to Internet protocols, but risks were low, since the NSFNET was not yet the critical infrastructure that the Internet would become. The NSFNET was created and administered by a tightly-knit research community characterized by strong interpersonal trust relationships. Individuals and specialized groups of this community - many of whom were involved with the NSFNET's predecessor, the ARPANET - took on functional roles for the development of Internet protocol standards, and for the allocation of critical Internet resources. The NSFNET was topologically laid out as a hierarchy, with the NSFNET backbone at the top, providing connectivity between regional network providers, which in turn provided connectivity to research institutions. The hierarchical topology of the NSFNET allowed for centralized control of inter-domain routing at the NSFNET backbone.

The nature of the risks and uncertainties inherent in the design of BGP are a consequence of the factors which shaped this environment: low risks, a small community with high trust relationships, and hierarchical control via topology. These factors would all change when the NSFNET was privatized to make way for the Internet. However, the basic form of BGP remained, as did the social formations and practices which ordered inter-domain routing. The distributed nature of the Internet was not by any means an inevitable technological outcome, but was rather produced from a context which was itself subject to hierarchical control.

Research on trust cautions that cooperation observed in the presence of an assurance structure should not be construed as a trust relationship. If the assurance structure fails, then it is likely that coordination will fail as well, in the absence of trust relationships which could provide the necessary warrants for coordination. An assurance structure is often interpreted

to be some kind of centralized institution, such as a government agency which warrants the value of money, or a legal institution which enforces contracts (Hardin 2002:195), or may even be a structure of incentives which shape cooperative behavior (Yamagishi and Yamagishi 1994). As my research shows, assurance structures may also arise from *topological* position within a network, when specific nodes - such as the NSFNET backbone - become gatekeepers by virtue of the power which accrues from their central position.

Although the NSFNET backbone acted as an assurance structure for the uncertainties inherent in BGP, it did not undermine the strong trust relationships amongst the technical communities developing and operating the NSFNET. These trust relationships predated the NSFNET itself, and were formed out of the collegial atmosphere of research communities, who needed to be tolerant of the risks and uncertainties attendant with failures in order to continue to develop the technologies of the Internet. The norms of openness, coordination and collaboration formed in this social environment formed the basis for the technical communities and institutions which emerged over the transition from the NSFNET to the Internet. As this account indicates, assurance structures may be an *outcome* of trust relationships, just as much as they are an *anchor* for cooperation from which trust relationships may emerge.

8.1.2 Trust in Governance

The nature of uncertainties in inter-domain routing are a consequence of the design decisions in BGP. These have remained relatively stable over time, as has BGP itself.³ However, the risks in inter-domain routing changed quite drastically over the transition from the NSFNET to the Internet. This was in part due to the changing perceived value of the Internet, as it went from being a research network to being critical infrastructure used by a wide variety of actors. These perceptions of value in turn shaped the risks and uncertainties taken on by those responsible for operating Internet infrastructure. Both risks and uncertainties changed due to the shift from the hierarchical topology - and centralized control - of the NSFNET, to a distributed topology of autonomous systems, with no central point of control.

Risk and uncertainty are shaped by topological position, as much as they are by the value attributed to a stable, reliable Internet. A more central topological position increases the risk taken on by an autonomous system, as the stability of the inter-domain routing system can be affected to a greater extent by the behavior of more central autonomous systems. Uncertainty is a consequence of both relative centrality, as well as of the number of directly interconnected autonomous systems. Both of these factors imply a greater number of possible points from which failures may occur, whether they are in the immediate neighborhood of an autonomous system, or if they flow from autonomous systems at a remove.

The distributed topology of the Internet's inter-domain routing system is stabilized by a distributed network of trust relationships, which cut across organizational boundaries, fol-

³There have, of course, been several extensions and additions to BGP, although none which have changed the basic nature of the protocol have found their way into production networks.

lowing the topology of interconnections amongst autonomous systems in the inter-domain routing system. The risks and uncertainties inherent in inter-domain routing are mitigated by the inter-organizational coordination and collaboration which these trust relationships enable. These trust relationships are produced and reproduced through the practice of managing inter-domain routing, and through the technical communities which sustain this practice, institutionalized as regional network operators' groups. At their most productive, trust relationships may allow for the relatively high degree of coordination and collaboration required to form an IXP, which in turn can itself then act as a productive site for trust relationships. These technical communities, and community-driven IXPs, are actively supported by more established technical communities (such as NANOG and RIPE), by non-profits (such as PCH and NSRC), and by sponsorship from states and for-profit entities. This degree of support, and linkage between technical communities, illustrates the importance of regional community to the development of the common practices and social relationships required to order the inter-domain routing system.

However, all communities are not created equal, nor do they always play similar roles, even though they are connected in a single global system of relationships required for the practice of inter-domain routing. The differences between NANOG and SANOG illustrate how the topological positions of autonomous systems within a region (internally, and relative to global topology), the nature of state involvement and regulation, and the relationships between states in a region, can all have effects on the formation of technical community. SANOG has a smaller core of personnel with strong trust relationships than NANOG, in part because of the relatively peripheral position of most autonomous systems in South Asia. As a result, trust relationships are not as salient to inter-domain routing as they are in North America, although they do contribute to the formation of IXPs in some South Asian countries.

The case of India shows how a highly regulatory and interventionist state, coupled with an oligopolistic market structure, can lead to conditions in which technical community - and accompanying trust relations - is only weakly formed. Most autonomous systems in India are doubly peripheral: peripheral to the global inter-domain routing system, and peripheral within India. Relationships from these peripheral positions are largely of confidence, rather than trust, in the oligopoly of large autonomous systems which provide connectivity into the global inter-domain routing system. While large autonomous systems in India need to maintain trust relationships with the technical communities around the world (maintaining a presence at venues such as NANOG) for operational purposes, peripheral Indian autonomous systems have no immediate need for such relationships with larger autonomous systems or amongst themselves, resulting in a relatively weak technical community in India. The consequence is a need for the state to provide outcomes (such as an IXP) which might otherwise be provided by a strong technical community.

The case of India demonstrates at once the success and failure of distributed governance: similar outcomes may be produced by means other than trust relationships and technical community within a specific region, under specific conditions. The principles under which these local articulations of governance operate may be quite different from those at play in

other regions or in global institutions. For instance, Internet governance activities in India are not quite as open and transparent as those at global Internet governance institutions. If this situation were common, it is likely that the governance of the inter-domain routing system would be substantially different from that which I have described, probably more centralized than distributed, and possibly managed under an international treaty organization, such as the ITU. Distributed governance allows for substantial variation in governance arrangements within its organization, as long as local articulations of governance can be reconciled with global ideals of governance articulated by more central technical communities (such as NANOG), and centralized governance institutions (such as the IETF, and the ICANN regime). Where topology is concerned, this reconciliation requires a compatibility of protocol, and associated social relations and practices, in the management of interconnection. Where resources are concerned, this reconciliation requires a formal relationship between centralized governance institutions, and local articulations of governance, in which local governance institutions agree to abide by the rules of the hierarchy of institutions above them (as in the case of IRINN conforming to APNIC policies).

At a global level, trust relationships are anchored by centralized institutions which serve specialized functions: the IETF for Internet protocol standards, and the ICANN regime for critical Internet resources. These centralized functions are, to varying degrees, made accountable to the Internet's technical communities through the representation of technical interests in these institutions, alongside political and economic interests. The mix, and relative power, of these interests varies depending upon the function of an institution.⁴ These assurance structures do not substitute for trust relationships, but rather fulfill specialized roles which are best carried out by centralized institutions. This is a key feature of distributed governance: specific functions are centralized only when it is not be practical to execute them in a decentralized manner.

The preference for trust relationships over assurance structures is as much the result of a specific historical trajectory, and technological form, as it is the articulation of a specific set of ideas about distributed networking. It is not for nothing that BGP uses the term “autonomous system” to describe the *autonomous* entities which interconnect to form the global Internet. Distributed governance functions through linked principles of interdependence and autonomy, allowing entities freedom to operate as they will within their own domains, while relying on trust relations and specialized centralized institutions to stitch these domains into a global whole.

8.1.3 Trust and Embeddedness

The stitching together of autonomous entities into a global whole does not, of course, occur in a vacuum. Openness and transparency, coordination and collaboration, must contend

⁴Networking equipment vendors such as Cisco, for instance, may be more concerned with participating in the IETF, than at ARIN. The ways in which the value of resources being managed - names vs. numbers - rubs up against intellectual property rights regimes, or scarcity, also shape the interests at play within an institution.

with moves towards closedness, competition and control arising from economic and political actors. Risk and uncertainty in inter-domain routing provide part of the justification for trust relationships which can cut across the organizational boundaries formed by economic and political interests. However, trust relationships are more than just a response to risk and uncertainty. They are also the means through which technical communities maintain the embeddedness of the markets and centralized institutions involved in producing the inter-domain routing system. In doing so, technical communities produce themselves as actors able to act “for the good of the Internet”, in relation to political and economic interests.

Just as the technical form of BGP encodes risk and uncertainty, calling for trust relations, it also encodes openness and transparency, providing the means through which autonomous systems may discover technical information about one another via routing announcements. Additionally, a number of projects exist which monitor network performance between autonomous systems, providing network administrators with the tools to more easily diagnose and fix the interconnections with other autonomous systems. The openness and transparency provided by these tools, and by BGP itself, provide a material basis for trust relations in the everyday practice of network administration, forcing the flow of information across economic and political boundaries. This creates a material basis for embeddedness - which I term “architectural embeddedness” - that technical communities utilize to form and connect to one another globally.

Conversely, the embeddedness of markets for interconnection is maintained through closedness, or rather, an openness restricted to the community of peering coordinators. In order to be able to negotiate peering agreements, peering coordinators must be able to share information with one another which would ordinarily be considered corporate secrets. They are able to do so because of the trust relations that they establish with one another. As I have found, a peering coordinator’s reputation and network of trust relationships are valuable in themselves: ISPs will hire peering coordinators who are respected and well-connected, anticipating that they will better be able to negotiate peering agreements. The market for interconnection of autonomous systems is embedded in a closed system of interpersonal relationships amongst peering coordinators, spanning corporate boundaries. As Granovetter (1985) points out, interpersonal trust relationships and reputation are essential mechanisms for the production of functioning, socially embedded markets. This is especially a concern where uncertainties over the quality of the commodity being traded are high (Kollock 1994), such as in the interconnection market, which requires ongoing monitoring of the stability and capacity of interconnections between autonomous systems.

Evans (1995) analyzes bureaucracies in terms of the degree to which they function autonomously from society, and the degree to which they are embedded in relations with society. He uses the term “embedded autonomy” to describe the condition in which bureaucracies are able to balance embeddedness and autonomy. A similar analysis may be made of the Internet’s centralized governance institutions, which must function autonomously to serve the common good, but must also be embedded in relations with the Internet’s technical communities (and society more generally) in order to gather information for their functions, and implement their functions. The major difference from Evans’ analysis lies in the fact

that the “professional bureaucracy” of the Internet is a mix of volunteer representation from technical communities and other actors (for the development of policies and standards), along with some professional staff to manage everyday processes (such as allocating resources, running conferences and maintaining websites).

In order to achieve autonomy, volunteers must be able to leave their organizational affiliations behind when engaged in governance activity.⁵ The principles by which these centralized institutions are operated, the specific nature of their work, and the representation of technical interests in relation to political and economic interests, shapes the nature of embedded autonomy. Some institutions, such as ARIN, are highly embedded in the technical community, and are also highly autonomous. For instance, ARIN has effectively repulsed moves by economic interests to treat IP address space as property to be traded, rather than as resources to be stewarded in service of the ARIN community. Others, such as ICANN, are dominated by economic interests, less embedded in the technical community, and less autonomous, in no small part because of the intrinsic value of the naming resources over which it asserts global control, positioning ICANN as a market maker for domain names.

As the case of India shows, architectural embeddedness, the embeddedness of markets, and the embedded autonomy of centralized governance institutions are not inevitable outcomes. Under conditions in which a strong technical community cannot form, markets may move towards disembeddedness,⁶ and centralized governance institutions may become overly embedded in economic and political relations, losing the ability to function autonomously. Embeddedness is subject to the same limits as the trust relations required to sustain it. Regardless of failures of embeddedness in local contexts, distributed governance continues to remain viable as long as these local contexts remain subject to the embeddedness of technical communities at regional and global scales. Recurrent struggles over Internet governance may be understood in terms of struggles over the embeddedness of technical communities across local, regional and global scales.

Embeddedness is not simply an analytical position, which shows that markets are always embedded in social relations. It also implies a political position, which - following Polanyi (2001) - argues that the attempt to disembed markets from society will result in social upheaval, as society is forced to conform to market logic, instead of markets serving the needs of society. Similarly, Evans (1995) states a political position, arguing that embedded autonomy of bureaucracy is required to enable a developmental state. The embeddedness of the markets and the centralized Internet governance institutions involved in inter-domain routing is socially valuable insofar as the Internet’s technical communities are able to act in

⁵The IETF operates upon the principle that individuals - not organizations - are participants in the Internet standards process. At ARIN meetings, individuals will speak in multiple roles, stating different positions as employees of companies, as ARIN AC members, or as ARIN community members.

⁶Here, I mean disembeddedness in the sense of Polanyi (2001), as a disembeddedness from social interests represented by technical community, insofar as they see themselves as acting in service of “the good of the Internet”. Following Granovetter (1985), markets are always embedded in social relations; in this case, it is just that social relations representing the economic interests of corporations come to dominate markets for interconnection.

service of a common ideal, “for the good of the Internet”. It is this position which makes the Internet’s technical communities at once *viable* and *meaningful* actors in relation to economic and political interests in inter-domain routing, constructing a form of private power which cannot be reduced to profit-seeking activity, functioning in service of the common good of a stable and reliable global Internet.⁷

8.1.4 Trust in Security?

The IETF’s Secure Inter-Domain Routing Working Group (SIDR WG) has approached the problems of risk and uncertainty in inter-domain routing by adding security features to BGP in the form of the Resource Public Key Infrastructure (RPKI) family of standards. RPKI functions through a “trust anchor” model of security: all RPKI-secured BGP announcements would be cryptographically signed by a root authority, allowing autonomous systems to establish the veracity of the claims made in a BGP announcement by verifying them with the root authority. RPKI goes a long way to reducing the risks and uncertainties in the inter-domain routing system, but does so at the cost of also reducing the “autonomy” of autonomous systems, since the RPKI root authority would effectively function as a policing agency for the inter-domain routing system. Such an agency would be able to exert unilateral control over the inter-domain routing system, a control which could potentially be leveraged political and economic interests to revoke routes to particular destinations on the Internet.

RPKI illustrates the way in which technical architecture can condition possibilities for governance, transforming the inter-domain routing system from an environment dependent upon distributed trust relationships, to one reliant upon a centralized institution managing the technical infrastructure of a trust anchor. By reducing the salience of trust relationships in inter-domain routing, such a technical architecture disintermediates architectural embeddedness. Unlike the functional specialization of centralized institutions that I have sketched thus far, this centralized institution would have to play the role of both RIR, and RPKI trust anchor, for resources assigned by the RIR. This is not idle speculation: all five RIRs already offer RPKI services to varying degrees, although adoption by autonomous systems has been limited. Issues of control and autonomy in RPKI have been hotly debated on the NANOG email list.

Cooperation in such a system relies on the stability of the assurance structure (the trust anchor), but as Hardin (2002:11) points out, this cooperation should by no means be mistaken for trust. Failure of the assurance structure would likely result in failure of the system as a whole, as there would no longer be any basis for cooperation to continue, in the absence of trust relationships. Nissenbaum (2004) offers a similar caution about “trustworthy” computing infrastructures, contending that trustworthy computing systems are better char-

⁷Where “stable” and “reliable” imply social expectations of freedom from surveillance and censorship, as much as the imply technological expectations of Internet infrastructure. For instance, following Edward Snowden’s revelations about the NSA’s activities, the IETF explicitly focused their standards efforts on supporting privacy and security to defeat surveillance; see <http://www.ietf.org/media/2013-11-07-internet-privacy-and-security.html>, last retrieved Aug 9th 2014.

acterized as providing surety, rather than trust. Nissenbaum's concern is that attempts to enable trust in online settings through more secure computing technologies will create a system that relies upon controlled access, fixed identities and perpetual surveillance. Such a system provides no guarantee of trust. It remains susceptible to malicious insiders, while potentially raising questions amongst ordinary users as to why their activities should be so extensively tracked, paradoxically reducing confidence in the system.

A secure technological solution, such as RPKI, which relies on a centralized assurance structure offers measurable, predictable security. However, it is relatively brittle, and subject to global failure, through a failure of technical infrastructure, and through capture by economic and political interests. It also reduces the possibilities for embeddedness of technical communities, as it diminishes the salience of the trust relationships to Internet governance.

In contrast, a distributed solution - such as the commonly used version of BGP - shifts the responsibility for risk and uncertainty to a network of trust relations. Such a system is subject to ongoing local failures, but rarely to global failure. Risks and uncertainties, then, are not merely undesirable factors to be engineered away, but must themselves be understood as being productive of trust relationships (Mathew and Cheshire 2010). The resultant form of distributed governance requires substantially more work to stabilize and order, but is resilient in the face of failures, and allows for greater local variation on global structures and practices of governance.

In making this argument, I am not speaking against creating more secure computing systems. Rather, I am pointing out that the problem of creating secure computing systems - and indeed, the very concept of what we mean by "secure" - requires attention to issues of both technology and governance. As I have shown, technology and the social formations required to govern it evolve alongside one another, and shape one another. They must be treated together as a unified design problem for large-scale infrastructures.

Clark et al. (2002) make a similar argument, advising protocol designers to be aware of "tussle spaces": areas of a design which may shape - and be subject to - the tussles between the many actors and interests involved in Internet infrastructure. As they point out, designers "are designing a playing field, not the outcome" (Clark et al. 2002:4). Their argument is that protocols should be designed in anticipation of tussle, with an explicit attempt made to separate the potentially tussle-free elements of a protocol from those which may be subject to tussle. As they note, this is by no means a straightforward or predictable process. I agree with their basic argument, but differ in my focus on outcomes. I believe that it is not sufficient - or even possible - to try to design protocols for tussle in themselves. Rather, I suggest that protocols and governance arrangements must be designed alongside one another; we must design these together, taking into account the possible range of outcomes that may arise from particular design choices. This is especially important in the context of the Internet, since the deployment of new technologies - and associated governance arrangements - must take into account the obduracy of existing technologies and governance arrangements. I have addressed these problems only in a limited fashion in this dissertation. However, I believe that the examination of patterns in technology and governance arrangements is critical to the future of the Internet, and to understanding how that future might shape the social,

political, and economic possibilities that the Internet provides for society.

8.2 The Production of Virtual Space

As ordinary users discovered the Internet in the 1990s, it came to be imagined as providing the means to escape the power of the nation state. John Perry Barlow captured and championed this image in his famous “A Declaration of the Independence of Cyberspace” (Barlow 1996), which opens:

Governments of the Industrial World, you weary giants of flesh and steel, I come from Cyberspace, the new home of Mind. On behalf of the future, I ask you of the past to leave us alone. You are not welcome among us. You have no sovereignty where we gather.

We have no elected government, nor are we likely to have one, so I address you with no greater authority than that with which liberty itself always speaks. I declare the global social space we are building to be naturally independent of the tyrannies you seek to impose on us. You have no moral right to rule us nor do you possess any methods of enforcement we have true reason to fear.

Governments derive their just powers from the consent of the governed. You have neither solicited nor received ours. We did not invite you. You do not know us, nor do you know our world. Cyberspace does not lie within your borders. Do not think that you can build it, as though it were a public construction project. You cannot. It is an act of nature and it grows itself through our collective actions.

Barlow’s articulation of the independence of virtual space - “cyberspace” - from territorial/physical space remains a powerful, enduring ideal of what the Internet is, and can be, in both academic and popular literature. I do believe that the Internet offers new possibilities for society, but it does not do so as an “act of nature”, but is rather actively produced in the ongoing efforts and struggles of the social formations and technologies involved in the distributed governance of Internet infrastructure. Virtual space is not a naturally occurring or inevitable space, any more than any other space which human beings inhabit. All of these spaces must be understood in terms of their means of production, as Lefebvre (1991) argued.

Lefebvre’s framework for understanding the production of space consists of spatial practice, representations of space, and representational space. A society produces space through spatial practice; which is informed by representations of space used by the technocracy of scientists, planners and so on; and which in turn produces a representational space of symbols and images overlaying physical space (Lefebvre 1991:38). Although Lefebvre’s primary concern is with the processes of the production of space in everyday life, his approach to space is invaluable for thinking about the production of the technological spaces of Internet infrastructure.

Neither virtual space, nor physical space, are mere containers within which human activity unfolds. Rather, they are actively shaped by human activity, and shape human activity in turn. It is only in the production of space through infrastructure that human activity can become extensive across space (at regional and global scales), and intensive within space (in urban agglomerations).

One of my principal aims in this dissertation has been to uncover the internal logic of the production of the virtual space of the Internet, and the manner in which the production of virtual space opposes and reconciles itself with the production of the spaces of the nation state, and the spaces of capital. This analysis provides insight into the internal logic of the Internet, as well as insights into the functioning of the modern nation state and capital, by examining the ways in which these systems interact with one another. It is critical that we form a more accurate spatial imaginary of the Internet, for without one, we will fail to understand the form and balance of powers involved in operating the present, and determining the future, of the Internet.

My examination of the production of virtual space proceeds in three stages. I begin by reviewing the technical architecture which enables the existence of virtual space. I then proceed to develop an understanding of the role that the intimate spaces of meetings of network operators' groups, and Internet governance institutions, play in the production of virtual space. Finally, I examine the interplay between the production of virtual space, and the production of the territorial space of the nation state. Together, these different modes of analysis help paint a unified picture of the processes, conflicts, and resolutions, involved in the production of virtual space.

8.2.1 The Architecture of Virtual Space

The production of any kind of space involves regulation and planning, which take the form of common standards (whether building regulations or communication protocols), blueprints, architectural plans and so on. The production of virtual space involves an architectural model which is formally separated across multiple layers, which provide the successive abstractions required to provide the appearance of a placeless virtual space over physical telecommunications infrastructure. It is the assumptions built into this architectural model - proceeding from norms, ideals and accepted practices - which provide the material basis over which distributed governance is made possible. As Winner (1980) argued, artifacts have politics. Improved industrial processes and machines may supplant labor. Some technologies - such as nuclear power - may call for centralized control, while others - such as rooftop solar power - may reduce the salience of central authority.

I treat the architecture of virtual space as having three layers: the "physical layer" of telecommunications infrastructure, the "logical layer" of Internet infrastructure, and the "application layer" which provides the appearance of virtual space. The logical layer is the moment of mediation between physical and virtual space, the first point at which identifiers

linked to physical infrastructure⁸ are abstracted to an identifier which is not intrinsically linked to any particular physical network interface, the IP address. This makes the architecture of the logical layer critical to understanding the means through which the placeless appearance of virtual space is produced; and, accordingly, the means through which virtual space is imagined as independent of physical space.

The logical layer provides a common set of protocols⁹ for communications across disparate physical telecommunications infrastructures, and a globally unique set of identifiers¹⁰ to manage routing of data. These abstractions embed notions of autonomy (“autonomous systems” in BGP), and related ideals of trust, openness, transparency, coordination and collaboration.

As I have shown, the “trusting” form that BGP took followed the ideals of the research community in which it was developed, and the hierarchical topology of the NSFNET. Over time, networks of trust relationships, accepted practices of network administration, and centralized institutional structures, evolved to manage the risks and uncertainties attendant with operating BGP, forming the system of distributed governance which I have described. This system of governance is made possible by the architectural form of BGP, which provides the means through which the Internet’s logical layer extends over space. The architecture of the logical layer therefore encapsulates two distinct but related concepts: *standards* for interconnection, and the *topology* of the things (autonomous systems in BGP) which are interconnected.

The assumptions of natural monopoly and territorial power built into the command-and-control governance structures of physical telecommunications infrastructure can only be effectively challenged by an architectural model which embeds alternative assumptions of governance. In spanning disparate physical infrastructures to form the global Internet, the Internet’s logical layer is the site at which stable interfaces are formed with the physical layer, while at the same time, the distributed governance of the logical layer rubs up against the forms of governance of the physical layer. Infrastructure in itself acts as a vehicle for power and authority, offering the material and shared symbolic means through which power and authority may be extended over space and time.

8.2.2 The Intimate Spaces of Internet Infrastructure

Responding to Winner (1980), Joerges argues that “built spaces always represent control rights. . . . Only rarely and in the most trivial senses can one show that such constraints are coupled to building form. In this view, it is the processes by which authorizations are built, maintained, contested and changed which are at issue in any social study of built spaces and technology” (Joerges 1999:424). Technological form does not determine politics, but rather shapes the possibilities for politics, just as politics shape and give power to technological

⁸Such as Media Access Control (MAC) addresses which identify network interfaces.

⁹IP and routing protocols such as BGP, as well as transport protocols such as TCP and UDP.

¹⁰IP addresses and ASNs.

form. Consider a counter-factual history: the problems of risk and uncertainty in inter-domain routing could have been solved through the creation of a centralized governance institution authorized to oversee and enforce routing configurations. Distributed governance does not spring full-formed from technical architecture alone, but rather must be infused with power and authority through social activity.

In my analysis, this social activity is split across three related sites. First, in the everyday practice of inter-domain routing, and the social relations of trust required to sustain this practice. Second, in network operators groups (such as NANOG and SANOG), which provide the means through which everyday practices are taught and refined, and through which social relations of trust are produced and reproduced over time and space. Third, in centralized governance institutions (such as the IETF, ICANN and ARIN), which provide the anchors for the social relations and everyday practices of inter-domain routing. The first of these - the practice and social relations of inter-domain routing - is only made possible by the work of network operators groups and centralized governance institutions.

Both network operators groups and centralized governance institutions use online tools, such as email lists, in their activities. However, as I have shown, it is only through regular meetings in shared physical space that these organizations can truly fulfill their missions. Boden and Molotch (1994) call this need for physical interaction “the compulsion of proximity”, in their discussion of the continued importance of “copresence” - in meetings of professional organizations, encounters at private clubs, and in other settings - to processes of globalization. They make two significant arguments which are of particular salience to my discussion here. First, they argue that the thickness of copresent interaction enables open-ended engagements and a greater density of shared information, which are required to deal with the complexity of the global systems - technological, economic and otherwise - of modernity. Second, they point out that copresence enables participants to form understandings with one another, without any record of their interaction. Copresence enables private interaction, at the same time as it provides the context for the transfer of tacit knowledge which cannot be encoded in explicit records of interaction (Polanyi 1966).

Copresence is integral to the formation of trust relationships, and the creation of a sense of generalized trust, at meetings of network operators groups. As I observed, and as my interviewees told me, the ability to look someone in the eye, to go out and have a beer with them, to share laughter in a room full of others with similar concerns, are all essential to the formation of trust, and community, amongst network administrators. Peering coordinators rely on copresence to construct spaces which are open to participation, but closed to recording; to participate, one must be physically present. These spaces enable more open discussion, and the sharing of sensitive information, which might not be possible if interactions in these spaces were recorded and made publicly available. In this sense, copresence is a necessary condition for the embeddedness of network interconnection markets.

Copresence is also essential to extension of trust relations, shared norms and symbolic resources, and notions of accepted best practices in network administration, across space. Practices must be disembedded from their points of origin, imagined as global absolutes, and then re-embedded in remote contexts. Trust relations, shared norms, and symbolic resources

follow, and enable, the dissemination of practice, providing the means to integrate practices across disparate geographies.

Network operators groups are, by their very nature, regional associations; the communities represented by network operators groups are spatially bounded, functioning to form intensive systems of trust relations within a region. In order to extend distributed governance across space, individuals must be willing and able to travel between network operators groups (and other Internet governance institutions) to construct extensive systems of trust relations linking geographically disparate social formations. The individuals who form these connections across space carry with them notions of accepted best practices in network administration, imparting these implicitly in the process of forming trust relations, and explicitly, by teaching technical workshops or delivering talks at conferences.

Two understandings of spatial practice are apparent from this analysis. First, of the practices involved in operating BGP, which produce the space of the Internet's inter-domain routing system. Second, of the practices of disembedding, re-embedding and integrating the practices and social relations required to produce the space of the inter-domain routing system. It is through these spatial practices that the representations of space required for inter-domain routing are transferred across geographies, and through which these are re-embedded as representational spaces for the practices of operating BGP within a geography.

Spatial practice implies a spatial division of labor, albeit not a spatial division by specialization and class relations (Massey 1994b), but rather one in which the individuals perform similar labor in different geographies, in order to interconnect their autonomous systems. The notion of geography becomes problematic here, as this analysis may better be presented in terms of the topological positions of autonomous systems within the inter-domain routing system. Power relations in the spatial division of labor for inter-domain routing are conditioned by topological position of autonomous systems, and also by the reputation and relationships that individual network administrators manage to cultivate. This brings us to a third understanding of spatial practice, which replaces geography with topology, complementing those which I have already discussed: that which is involved in producing the appearance of placeless virtual space over physical infrastructure, and across geographies.

The essential social elements of distributed governance - trust, community, practice and embeddedness - are forged through copresent interactions in intimate spaces. Online interactions - such as on email lists - cannot supplant copresent interactions, but are invaluable to the ongoing maintenance of social relations between the times and spaces of copresent interaction. Paradoxically, the production of seemingly placeless virtual space depends upon the intimate spaces of physically co-located gatherings.

8.2.3 Where Virtual Space Meets Territorial Space

The spatial imaginary which attaches to distributed governance is that of a network, in which power must be understood in terms of topology. The architecture of the Internet's logical layer constructs an abstract network of autonomous systems and IP addresses, overlaying and interconnecting networks of physical telecommunications infrastructure. The social el-

ements of distributed governance construct multiple social networks, which are sometimes distributed - trust relations amongst network administrators, relationships between network operators groups - and sometimes embedded in centralized governance institutions.

In contrast, the territorial space of the nation state is conventionally thought of as physical space bounded by established borders. This is not a static “timeless” space, merely occupied by the state, and acting as a container of relationships (Agnew 1994), but is actively produced, “mapped, modified, transformed by the networks, circuits and flows that are established within it - roads, canals, railroads, commercial and financial circuits . . .” (Lefebvre 2003:84). These “spatial fixes” (Harvey 1985) of infrastructure are just as important to processes of globalization, as they are to the extension of the power of the state within its territory, contextually enabling and limiting these processes. These dynamics speak to the extension of the power of the state outside its territory, insofar as states can be regional powers (such as India in South Asia), or hegemonic powers within the world system.

The problem of the interaction between the virtual space of the Internet, and the territorial space of the state, can therefore be conceived in terms of overlapping logics of spatial production through infrastructure. While these are distinct logics, it is not that one causes the retreat of the other, but rather that these spaces interpenetrate one another. They must be understood in terms of the processes through which they manage the conflict and resolution of their distinct logics, as topologies of power (Allen 2011). Through these processes, the Internet is produced in the state (in terms of topology and institutional structures), just as the state is produced in the Internet (through regulation and integration into global systems of governance).

The evolution of the Internet in India illustrates these processes well. The introduction of the Internet to India involved struggles between different parts of the state: the Department of Telecommunications which controlled domestic telecommunications infrastructure, and the Department of Electronics. The initial struggles were over the development of a domestic computer network to link education and research institutions; later struggles involved the provision of Internet connectivity to the nascent India software industry. Through these struggles, the authority of the Department of Telecommunications was gradually eroded, beginning with the creation of government-run corporations to manage telecommunications infrastructure, then with the provision of a new legal vehicle (the Software Technology Parks of India) to enable Internet connectivity for the software industry, and finally with the opening of the telecommunications sector (including the provision of ISPs) to private actors.

The history of the Internet in India is by no means simply an account of a one-way process of the erosion of state control. The state’s involvement in the creation and operation of the only IXP in India, NIXI, illustrates the manner in which the territoriality of the state reconciles itself with the Internet, as the state introduced itself into a key component of the Internet’s topology within India. The new capacities that the state created internally for the operation of NIXI contributed to integrations into the centralized institutions of Internet governance, with the creation of IRINN - under an agreement with APNIC - for the provision of critical Internet resources within India.

The processes of conflict and resolution between the production of territorial space, and

the production of virtual space, are not frictionless or perfect. For instance, the market power of large ISPs in India limits the utility of NIXI, thereby limiting the topological power of the state in the Internet. Similarly, while the creation of IRINN as a state agency is a point of pride for NIXI, IRINN can only function insofar as its policies are in consonance with that of the supranational private power of the ICANN regime, represented by APNIC. This is not to suggest that the state's power is diminished, but rather that it is changed. After all, the Internet can only enter territorial space *through* the state. The critical question to ask is to what degree the state is aware of, and involved in, the entry of the Internet.

NIXI, IRINN, and the regulatory regime for ISPs in India, represent the outcomes of processes of conflict and resolution between the territorial logic of the Indian state, and the logic of the production of the virtual space of the Internet. The Internet did not cause the state to wither away; rather, the state developed capacities to engage with the Internet, just as the Internet developed mechanisms for engaging with the state.¹¹ Through these processes, the state and the Internet enmesh themselves in one another's topological and institutional arrangements.

Similar processes occurred in the case of the USA, although these involved the state surrendering control of topology - by privatizing the NSFNET and funding the creation of IXPs - while at the same time providing the institutional backing for private power to assert control over critical Internet resources, through the creation of the ICANN regime.¹² As with the evolution of the Internet in India, this was not a straightforward process; it involved the formation of a tenuous bargain between private power, the Internet's technical community, the US state, and various other governments (Mueller 2002). As a hegemonic power promoting a free-enterprise system, the USA conceived of the power of private enterprise as a means to extending its own interests (Arrighi 1990). These private powers were well-established in the control of Internet topology, and in representation at the IETF. As Arrighi points out, the free-enterprise ideology of US hegemony created the conditions under which all states, including the USA, and the broader inter-state system, are no longer the primary locus of power in the world system:

... the contemporaneous development of international organizations and transnational corporations has created an extensive and dense network of pecuniary and nonpecuniary exchanges which no single state can control unilaterally and, more importantly, from which no state can "delink" except at exorbitant costs (Arrighi 1990:403).

¹¹Through the development of APNIC policies for the provision of NIRs such as IRINN. Engagements between the logic of the Internet and the logic of states go beyond just NIR policies, of course, extending to include the formal representation of states in ICANN's multistakeholder model, discussions in the UN-sponsored Internet Governance Forum, and recurrent conflicts over the role of state authority in Internet governance activity.

¹²Recall that ICANN is incorporated as a Californian non-profit entity, and was created under an agreement with the US Department of Commerce.

What is distinctive about the Internet is that the nature of private power cannot be reduced to that expressed by economic interests. As I have shown, the Internet's technical communities, practices and technology construct a different kind of private power, which acts through distributed governance, "for the good of the Internet". The production of virtual space occurs in the interaction between the various interests of state, market and Internet.

8.3 Distributed Governance for Internet Infrastructure

8.3.1 Towards a Model of Distributed Governance

The distributed governance of the Internet's inter-domain routing system works through shared practices, dependent upon trust relations, produced and reproduced in technical communities, and anchored by centralized governance institutions. These exist in relation to technological forms which embed social ideals of autonomy, openness, transparency, coordination and collaboration. These ideals in turn inform the production of practices and relations, as well as of centralized governance institutions constructed to perform specialized functions which anchor the practices and relations of distributed governance. This system reproduces itself over space through technological form, and through regional technical communities which act as channels for the transmission and refinement of practices and ideals, and hubs for the formation of trust relations.

In this account, risks and uncertainties in technological form are a feature which enable autonomy. Trust relations - and associated practices of openness, transparency, coordination and collaboration - are required to manage these risks and uncertainties. To completely engineer risk and uncertainty away would be to shift the range of possibilities for governance away from distributed governance, and towards centralized authority and policing.

Technological form is, therefore, of significant importance to the construction of distributed governance. Power shapes, is shaped by, and extends spatially through technological form. However, power does not reside within technological form; it flows from the system of trust relations and practices of geographically distributed technical communities.

Distributed governance depends upon the embeddedness of technical communities in order to function. It reaches its limits in the absence of a strong technical community; in which case state and market actors may step in to provide goods which might otherwise be provided through coordination and collaboration amongst technical communities. While the powers of state agencies, and market actors may become dominant under certain local conditions, they must act in accordance with the logic of distributed governance at a global scale. Localized segments of the inter-domain routing system's topology may be organized under their own distinctive logic, but to integrate with the overall topology of the inter-domain routing system, they must reconcile themselves with the logic of distributed governance.

The related elements of technological form (standards and topology) and of social form (trust relations, practices, technical communities, centralized governance institutions and

embeddedness), as they extend over space, together constitute the “architecture of trust” of distributed governance. This constructs distributed governance as an internally coherent system, and defines the ways in which it engages with, shapes, and is shaped by, other systems of governance, whether state or market.

8.3.2 Perspectives on Internet Governance

Internet governance encompasses a much wider set of concerns than those which I have dealt with in this dissertation, including issues of privacy, surveillance, security, intellectual property rights, access to the Internet, and more. My focus is on a specific component of Internet infrastructure - the inter-domain routing system - which enables the very existence of the Internet as an interconnected set of networks. This is perhaps the most critical component of Internet infrastructure, in that it provides the basis over which the broader set of concerns and governance issues relating to the Internet are made possible. In examining this system, I formulated the concept of distributed governance, which offers a viewpoint distinct from other perspectives on Internet governance, of which I survey a representative set here.

Goldsmith and Wu (2006) argue that nation states are employing coercive power to shape the global Internet, and to reconstruct their borders in cyberspace. They are careful to say that governmental coercion is not the only factor of importance to Internet governance. Regardless, they state strongly that “the failure to understand the many faces and facets of territorial governmental coercion is fatal to globalization theory as understood today, and central to understanding the future of the Internet” (Goldsmith and Wu 2006:184).

In contrast, Foster and McChesney (2011) insist on the primacy of capital, viewing governments as accomplices in the production of monopolies across wide swathes of the Internet. Their analysis focuses in large part on the application layer. They believe in the democratic promise of the Internet, but argue that this can be attained only through widespread public resistance to the encroachments of capital.

Johnson et al. (2004) posit that the governance of the online world is substantially different from that of the offline physical world. Drawing on Benkler (2007), they argue that Internet governance may be produced through the aggregated actions of individuals in “peer production”:

Rather than electing representatives, or hoping that some unaccountable online sovereign will do the right thing, we can collectively assume the task of making and implementing the rules that govern online activity ... The aggregation of numerous individual decisions about who to trust and who to avoid will create a diverse set of rules that most accurately and fairly serves the interests of those who use the online world (Johnson et al. 2004:6-7).

DeNardis and Raymond (2013) offer a critique of multi-stakeholder governance, through the creation of a typology of governance types. They caution against the common notions

that all Internet governance is multi-stakeholder governance, and that multistakeholder governance is a single form of governance, which in itself serves the public interest. For instance, they analyze network interconnection for inter-domain routing as being constructed through administrative decisions and contracts in the private sector, which they point out cannot be understood in terms of multi-stakeholder governance. Their typology of governance is based upon four stakeholder types - states, inter-governmental organizations (IGOs), corporations and non-governmental organizations (NGOs) - and three kinds of authority relations - hierarchy, polyarchy and anarchy. In their account, different kinds of arrangements of stakeholders, authority relations and procedural rules can lead to different kinds of governance, which may to different degrees be multi-stakeholder in nature.

Mueller (2010) develops the concept of Internet governance as “networked governance”, building on theories of network organization which originate with Powell (1990). He shows how the Internet is governed through transnational, networked relations, which require actors who wish to participate in Internet governance to engage with these networks in meaningful ways. His principal concern is with the ways in which states engage with networked governance, which may range from conflict between state hierarchies and governance networks, to the integration of states into governance networks. In more recent work, Mueller et al. (2013) extend the concept of networked governance to understand the mechanisms used to secure the inter-domain routing system.

8.3.3 The Practice of Infrastructure

Each of the perspectives on Internet governance surveyed in the last section share certain common features, even though they approach the problem of Internet governance quite differently. First, they are all structural explanations; they cannot adequately account for processes of genesis and change of the structures which they describe. Second, they tend to favor certain units of organization over others - whether individuals, states, corporations, IGOs or NGOs - to describe how these units are structured in Internet governance. Finally, they make claims about which units of organization do, or should, hold power in Internet governance. Mueller (2010), and DeNardis and Raymond (2013) are exceptions to this last point, since their focus is on how different units of organization may together constitute a single system of governance, through networked or multi-stakeholder arrangements.

My work diverges from these perspectives on Internet governance in two distinct ways, to contribute to a more complete theory of Internet governance. First, in my explicit focus on the the *practice* of inter-domain routing. Second, in my examination of the mechanisms required to sustain this practice: *community* and *trust*. My explanation of distributed governance therefore begins with an understanding of the practice of infrastructure, and thereby illustrates how practice - and associated social and technological formations - provides the basis for governance within a distributed system.

Through this analysis, I reassert the importance of the Internet’s technical communities to Internet governance, treating them as an analytical category on par with states and corporations. Other perspectives on Internet governance do take account of technical com-

munities, but typically deal with them either as a historical artifact, or fail to fully develop a theory of their role in Internet governance. The Internet's technical communities cannot be treated simply as agents of state or corporate power; members of these communities must be able to reconcile the obligations and ideals they hold with their principals with those which they hold in their communities. Technical communities must be understood as constitutive of governance through the ideals they hold, the norms under which they operate, the practices they form and propagate, the role of trust relations in the technologies they administer, and the manner in which they maintain the embeddedness of the markets and centralized governance institutions of Internet infrastructure.

The distributed technological arrangements of Internet infrastructure call for distributed social arrangements of governance. Political power is located in Internet infrastructure in these distributed social arrangements, and in their interaction with the power of states and corporations. It is these distributed social arrangements of governance which make the Internet unique in comparison to other infrastructures, and which allow us to imagine and enact possibilities for freedom in online settings beyond those which we experience in the physical world.

8.4 Infrastructural Power in the Information Society

The concept of the information society represents a vision of a world in which information structures dominant processes in society. This is a developmental move from industrial society, which in turn was a development over agricultural society. Each kind of society is characterized by a particular mode of production, which shapes the nature of social relations, power and authority.

Major theorists have been trying to make sense of the form of society which comes after industrial society at least since the 1970s, proposing a range of concepts, from "post-industrial society" (Bell 1973), to "post-modern society" (Harvey 1990, drawing from Lyotard 1985), to "network society" (Castells 2000). Not all of these theories take information as their primary problematic, but they do trace certain common features: an increased mobility of capital, intensifying processes of globalization, and the rise of the services economy. These features point to a collapse of space and time, although different theorists treat this phenomenon differently.

Some theorists proceed from the assumption that space and time can never be completely collapsed, and are therefore most concerned with the mechanisms through which the experience of a collapsed space-time is produced. Harvey (1990) develops the concept of "time-space compression", to explain the means through which capital engages in spatial fixes for its problems (through movement in space), and increases the speed of circulation of commodities, both of which can only be enabled through spatial fixes of capital in infrastructure (for transportation, communications, power and so on). In contrast, Giddens (1991) is more concerned with the manner in which social relations can be extended over space and time, through "time-space distancing". He argues that the disembedding of

social relations from local contexts is made possible through trust in symbolic tokens (such as money) and in expert systems (the professional expertise of lawyers, doctors, architects and so on).¹³

Other theorists assume that space and time can be perfectly collapsed, allowing them to construct totalizing perspectives on society. In the “network society” of Castells (2000), the world is divided between a “space of flows” and a “space of places”, which cannot be reconciled with one another. In the “space of flows”, time and space are not merely compressed, but collapse to nothingness, structured by “light-speed-operating information technologies” (Castells 2000:501); while the “space of places” is constrained to the physical geography of fragmented localities. Barlow’s “A Declaration of the Independence of Cyberspace”, cited earlier, is representative of this perspective in popular culture, as is the oft-repeated comment that “information wants to be free”.¹⁴

I take strongest issue with this latter perspective, in its inability to account for difference, and its technologically determinist narrative of a uniform, undifferentiated end to which all of society must strive. Writing about the rhetoric of technology in futurism, Carey and Quirk (1988) make a similar case, arguing that such a perspective reflects “the desire to whittle down valuable forms of conduct and modes of life to a single set of consistent principles”, which “is only possible in a thoroughly technologized landscape where machines alone possess teleological insight” (Carey and Quirk 1988:200). Theories which imagine society as existing over perfectly collapsed time and space are representative of what I think of as a two-fold disembeddedness which is common in thinking about the information society: a disembeddedness of social relations from space and time, and disembeddedness of information from social relations, both of which are thought to be accomplished through technology.

All of the theories of the information society I have surveyed face a common shortcoming. In their focus on the experience of everyday life, they fail to pay attention to the spatial practices involved in the production of infrastructure. For Harvey, infrastructure is a spatial fix of capital; for Giddens, it is an expert system to be trusted; and for Castells, it is purely technology. To make sense of the modern world - of the information society - we must pay attention to the details of the production of space through infrastructure. As Massey (1994b:22) argues, “since social relations are bearers of power what is at issue is a geography of power relations in which spatial form is an important element in the constitution of power itself.”

In the case of the information society, the spatial form at stake is, by and large, the Internet. Through an examination of the distributed governance processes which order and

¹³Giddens’ use of trust is different from the sense in which I have used it, and is more akin to confidence, or faith. See Chapter 2 for a more complete discussion of these differences.

¹⁴Originally stated by Stewart Brand at the first Hacker’s Conference in 1984. The complete version of his statement is somewhat more nuanced: “On the one hand information wants to be expensive, because it’s so valuable. The right information in the right place just changes your life. On the other hand, information wants to be free, because the cost of getting it out is getting lower and lower all the time. So you have these two fighting against each other.” See <http://www.rogerclarke.com/II/IWtbF.html>, last retrieved Jul 17, 2014.

stabilize Internet infrastructure, I have tried to contribute to a theory of infrastructural power in the information society. Mann (1984) develops a theory of infrastructural power in the context of the state, arguing that infrastructural power - expressed through means such as literacy, coins, weights and measures, and communication and transport - represents the autonomous power of the state in its capacity to penetrate civil society. Yet, as he points out, infrastructures are not specific to the state; they are part of general social development, and may be exploited by actors other than the state, such as corporations.

The Internet is the most recent of these infrastructures of general social development. Like Mann's infrastructures, the Internet penetrates society; and as the principal infrastructure of the information society, offers significant power to those who control it. As I have shown, the infrastructural power of the Internet cannot be understood as being "controlled" in the conventional sense, but must instead be understood in terms of what I have called distributed governance. The infrastructural power of the Internet accrues to a complex socio-technical formation, in which technical communities hold autonomous power in relation to states and corporations.¹⁵

Infrastructural power in the Internet - and thereby in the information society - must be understood in relational terms. It is structured through relations amongst technical communities, states and corporations; by relations within and between technical communities; and through the information shared over these relations.

In this sense, information is not an object to be passed around, or possessed of a will of its own, wanting to be free. Information is a relational quantity, which only exists over social relations, and makes these social relations possible. Information is not disembedded from social relations, nor are social relations disembedded from space, in the information society. Rather, their embeddedness is constitutive of the production of the space of the information society, produced through layers of infrastructure. So where in the world is the Internet? It is in social relations embedded in, and productive of, space; in information embedded in, and productive of, social relations; with both of these conditioned, but not determined, by technology. If we are able to ascribe ideals of "freedom" and "democracy" to the Internet, it is not because of any immanent quality of technology, but rather because of the arrangements and quality of the socio-technical relations that constitute the distributed governance of the Internet.

¹⁵With notable exceptions such as Bell (1973), the constitution of modernity is often considered primarily in terms of the interplay between states and corporations, with little attention to technical communities. This is an oversight which I seek to remedy here.

Afterword

As I was writing this dissertation, a number of events unfolded which brought greater public attention to Internet infrastructure, raising questions of where and how exactly political power is located in Internet infrastructure; questions I have tried to address throughout this dissertation.

Edward Snowden's revelations about the activities of the NSA and allied intelligence agencies raised issues of surveillance and privacy on the Internet, and more generally in global telecommunications infrastructure. The Brazilian government responded by committing to investments in telecommunications infrastructure which would bypass the USA.¹⁶ The German government proposed a stronger pan-European Union telecommunications infrastructure, which would keep traffic within the EU as much as possible, to frustrate access by the NSA and allied intelligence agencies.¹⁷

In July 2014, the UK government rushed through legislation, with little debate, which would force ISPs and telephone networks to indefinitely retain records of all data traffic and calls that they carry, even though such legislation runs counter to a ruling by the European Court of Justice.¹⁸ Over the last few years, the Russian government has gradually enforced greater controls over the Internet, most recently passing new laws which allow the government to block websites without explanation, which force bloggers to register with the Russian mass media regulator, and which require Internet companies to host information generated by Russian users on servers within Russia.¹⁹ China, of course, remains amongst the most controlled media environments in the world, although it must be noted that China maintains its so-called "Great Firewall" as much through carefully engineered topology with central points of control, as it does through legal and regulatory regimes.²⁰

In 2014 in the USA, the structure of peering agreements changed, as Netflix - a provider of streaming video services - was forced to pay Comcast, Verizon and other large American ISPs

¹⁶See <http://www.wired.co.uk/news/archive/2014-02/25/undersea-cable-brazil-europe>, last retrieved Aug 10th, 2014. Much of South America's international telecommunications links are currently channeled to international locations through the USA.

¹⁷See <http://www.bbc.com/news/world-europe-26210053>, last retrieved Aug 10th, 2014.

¹⁸See <https://www.openrightsgroup.org/campaigns/no-emergency-stop-the-data-retention-stitch-up>, last retrieved Aug 10th, 2014.

¹⁹See <http://www.bbc.com/news/technology-28583669>, last retrieved Aug 10th, 2014.

²⁰See <http://www.theatlantic.com/magazine/archive/2008/03/-the-connection-has-been-reset/306650/>, last retrieved Aug 10th, 2014.

for high speed access to Netflix subscribers in these ISPs.²¹ This raised issues of “network neutrality”, the principle that individual services should not have to pay ISPs to reach their users; nor should users have to pay ISPs for preferential access to individual services. In the interests of understanding these issues, and potentially regulating them, the US Federal Communications Commission (FCC), which regulates media and telecommunications in the USA, asked for, and received, the confidential peering agreements that Netflix entered into with these ISPs.²²

These events could be read as an assertion of state and corporate control over Internet infrastructure. However, it is important to consider the mechanisms through which these events proceeded. These events involved states exerting territorial, and extraterritorial power through topology, and through legal regimes; and corporations exerting control by virtue of topological position. In some cases, states sought to evade the surveillance capabilities of other states, as with Brazil and Germany against the NSA. In others, states sought to increase territorial control of information flows, as with the UK, Russia and China. Corporations running ISPs sought to leverage the power that accrues by virtue of their topological position, in their control of the “last mile” of access to users of the Internet. In addition, states, corporations and the Internet’s technical communities continued to tussle over control of the Internet’s centralized governance institutions, as I detailed in Chapter 5.

The responses to these events cannot help but take place over the terrain upon which they unfolded: the technology and governance arrangements of Internet infrastructure. Responses range from lobbying for new legal regimes (whether against surveillance, or for network neutrality), to re-asserting control over topology, to negotiating the balance of power in the Internet’s centralized governance institutions. In every instance, the Internet’s technical communities play a critical role, by crossing state and corporate boundaries for the continued coordination and collaboration required to operate the global Internet, and by lending their voices and expertise to shape political debate.

In 1986, Clifford Stoll, an astrophysicist-turned-network administrator at the Lawrence Berkeley National Lab noticed a discrepancy in the network he managed. This led to an investigation of what turned out to be one of the first instances of espionage on the nascent Internet. In the book in which he documented this investigation, he wrote at length about the coordination and collaboration, and the relationships of trust, which were needed to uncover those responsible for the espionage:

Networks aren’t made of printed circuits, but of people. . . My terminal is a door to countless, intricate pathways, leading to untold numbers of neighbors. Thousands of people trust each other enough to tie their systems together. Hundreds of thousands of people use those systems, never realizing the delicate networks that link their separate worlds. . . all those people work and play unaware of how

²¹See Chapter 4 for a discussion of peering agreements, and the issues that arise with CDNs and high bandwidth services such as streaming video.

²²See <http://arstechnica.com/tech-policy/2014/06/fcc-gets-comcast-verizon-to-reveal>, last retrieved Aug 10th, 2014. Network neutrality is a complex issue, which I will not examine in any detail here.

fragile and vulnerable their community is. It could be destroyed outright by a virus, or, worse, it could consume itself with mutual suspicion, tangle itself up in locks, security checkpoints, and surveillance, wither away by becoming so inaccessible and bureaucratic that nobody would want it anymore. . . . I don't want to be a computer cop. I don't want our networks to need cops (Stoll 2005:394-395).

As I have shown in this dissertation, the Internet we have today is not so different from that which Stoll inhabited. Networks are still tied together by relationships of trust; and it is these relationships of trust, in distributed governance, which allow people to attribute "freedom" and "democracy" to the Internet. My aim is to bring a fresh awareness of the importance and value of the Internet's technical communities to the governance of the Internet. Recent events threaten the emergence of an Internet which is less free and democratic, whether through ongoing surveillance activities, through balkanization into regional or national internets, or through increasing corporate control. I believe that the realities of interdependence and the value of a single global Internet make it unlikely that such a future will come to pass, even though it is possible that the Internet will become more controlled and monitored across the world. Like Stoll, I don't want our networks to need cops. My hope lies in the system of distributed governance which I have outlined in this dissertation, acting "for the good of the Internet".

Glossary

AfriNIC	Africa Network Information Center, the RIR for the African region.
ALAC	At-Large Advisory Council, the ICANN stakeholder group representing the interests of Internet users.
ANS	Advanced Network Services, the non-profit entity which took over operation of the NSFNET backbone.
APNIC	Asia-Pacific Network Information Center, the RIR for the Asia-Pacific region.
APRICOT	Asia-Pacific Regional Internet Conference on Operational Technologies
ARIN	American Registry for Internet Numbers, the RIR for the North American region.
ARIN AC	ARIN Advisory Council
ARPA	Advanced Research Projects Agency of the US Department of Defense.
ARPANET	A research network created by ARPA.
AS	Autonomous System
ASIC	Application-Specific Integrated Circuit
ASN	Autonomous System Number, the unique number used to identify an autonomous system.
ASO	Address Supporting Organization, the ICANN stakeholder group representing the interests of the RIRs.
BARRNET	Bay Area Regional Research Network, a regional network in the NSFNET.
BBN	Bolt, Beranek and Newman, the firm responsible for building and operating the ARPANET network layer.
BBS	Bulletin Board Service
BCP	Best Current Practice, the IETF document series capturing recommended practices in network administration.
BDIX	Bangladesh Internet Exchange
BGP	Border Gateway Protocol, the protocol which supplanted EGP as a means of interconnecting gateways on the NSFNET, and later the Internet.
BOF/BoF	Birds of a Feather, gatherings of people with similar interests at NANOG meetings.
BSD	Berkeley Software Distribution, a UNIX distribution originally managed by a group at the University of California, Berkeley.

CalREN	California Research and Education Network
CaribNOG	Caribbean Network Operators Group
ccNSO	country code Name Supporting Organization, the ICANN stakeholder group representing the interests of country code top-level domain operators.
ccTLD	country code Top-Level Domain
CCIE	Cisco Certified Internetwork Expert
CCITT	Comité Consultatif International Téléphonique et Télégraphique, an international standards organization for telecommunications, which was succeeded by the ITU-T.
CCNA	Cisco Certified Network Associate
CCNP	Cisco Certified Network Professional
C-DAC	Center for the Development of Advanced Computing, a computer science research center in India.
CDN	Content Delivery Network
CENIC	Corporation for Education Network Initiatives in California
CIDR	Classless Inter-Domain Routing, a mechanism for specifying IP address blocks with prefixes of arbitrary lengths.
CLNP	ConnectionLess Network Protocol, an OSI protocol which was proposed as a replacement for IP by the IAB in 1992.
CNRI	Corporation for National Research Initiatives
DEITY	Department of Electronics and Information Technology, a department in the Ministry of Communications and Information Technology of the Government of India.
DHS	Department of Homeland Security, a department of the US government tasked with internal security.
DNS	Domain Name System
DROP	Don't Route Or Peer, a list of BGP routes which should be blocked, maintained by Spamhaus.
EDROP	Extended Don't Route Or Peer, and extended version of the DROP list.
EFF	Electronic Frontier Foundation, a non-profit advocacy organization dedicated to protecting freedoms on the Internet.
EGP	Exterior Gateway Protocol, the protocol first used to interconnect gateways on the NSFNET.
EISPAI	Email and Internet Service Providers Association of India, the original name of ISPAI.
ERNET	Education and Research Network, a network connecting education and research institutions in India.
FBI	Federal Bureau of Investigation
FCC	Federal Communications Commission
FIB	Forwarding Information Base
FIPS	Federal Information Processing Standard

FLAG	Fiber Link Around the Globe, a submarine cable system.
GAC	Governmental Advisory Committee, the advisory committee through which governments are represented at ICANN.
GGP	Gateway to Gateway Protocol, the protocol used to interconnect gateways on the ARPANET.
GNSO	Generic Names Supporting Organization, the ICANN stakeholder group representing the interests of generic top-level domain operators.
GOSIP	Government OSI Profile, the version of OSI used for government procurement.
GPF	Global Peering Forum, a meeting series for peering coordinators.
gTLD	generic Top-Level Domain
HKIX	Hong Kong Internet Exchange
IAB	Internet Architecture Board; formerly Internet Activities Board, which was in turn a reformulation of the ICCB.
IAD	IETF Administrative Director
IANA	Internet Assigned Numbers Authority, the entity responsible for ensuring the uniqueness of assignments of Internet numbers such as IP addresses and well-known port numbers.
ICANN	Internet Corporation for Assigned Names and Numbers
ICCB	Internet Configuration Control Board, predecessor to the IAB.
IDR WG	Inter-Domain Routing Working Group, a working group at the IETF which focuses on the development of BGP and allied protocol standards.
IESG	Internet Engineering Steering Group
IETF	Internet Engineering Task Force
IGF	Internet Governance Forum
IGO	Inter-Governmental Organization
IGP	Interior Gateway Protocol, a class of routing protocols used within autonomous systems or closed networks.
IMP	Interface Message Processor, a specialized computer providing connectivity for hosts to the ARPANET.
INOC-DBA	Inter-Network Operations Center Dial-By-ASN
InterNIC	Internet Network Information Center
IP	Internet Protocol
IPTO	Information Processing Techniques Office, the ARPA office which oversaw the creation of the ARPANET.
IRC	Internet Relay Chat
IRINN	Indian Registry for Internet Names and Numbers
IRR	Internet Route Registry, a database of inter-domain routing information, such as RADb.
ISO	International Organization for Standards
ISOC	Internet Society, the organizational home for the IAB and the IETF.
ISP	Internet Service Provider

ISPAI	Internet Service Providers Association of India
ITRs	International Telecommunications Regulations
ITU	International Telecommunications Union
ITU-T	The telecommunications standardization sector of the ITU.
IXP	Internet Exchange Point
LACNIC	Latin America and Caribbean Network Information Center, the RIR for the Latin America and Caribbean region.
LAN	Local Area Network
LIR	Local Internet Registry
MAE-East	Metropolitan Area Ethernet East, an early IXP.
MAE-West	Metropolitan Area Ethernet West, an early IXP.
MED	Multi-Exit Discriminator, a BGP attribute used to control routing behavior between autonomous systems.
MENOG	Middle East Network Operators Group
MPLS	Multi-Protocol Label Switching
NACR	Network Add/Change Request
NAP	Network Access Point, NSF-funded locations at which networks to interconnect after the privatization of the NSFNET backbone network.
NANOG	North American Network Operators Group
NAT	Network Address Translation
NCP	Network Control Program, the software used by ARPANET IMPs to connect to one another.
NDA	Non-Disclosure Agreement
NGO	Non-Governmental Organization
NIR	National Internet Registry
NIXI	National Internet Exchange of India
NLNOG	Netherlands Network Operator Group
NLRI	Network Layer Reachability Information
NOC	Network Operations Center
NomCom	Nominating Committee
NPIX	Nepal Internet Exchange
NRO	Number Resources Organization
NRPM	Number Resources Policy Manual
NSFNET	The NSF-funded successor to the ARPANET.
NSRC	Network Startup Resource Center, a non-profit organization which supports the training of network administrators, and the development of the Internet around the world.
NTP	Network Time Protocol
NWG	Network Working Group, the group of researchers who developed the host layer protocols for the ARPANET.
OSI	Open Systems Interconnection, the networking standard created by the ISO.

OSPF	Open Shortest Path First, a popular IGP.
PA	Provider-aggregatable address space
PCH	Packet Clearing House, a non-profit organization which supports training of network administrators, and setting up of IXPs around the world, and provides a wide range of services to the worldwide network operations community.
PDP	Policy Development Process, the ARIN process for development of policy.
PGP	Pretty Good Privacy
PI	Provider-independent address space
PIE	Pakistan Internet Exchange
POC	Point of contact
PPC	Public Policy Consultation, an abbreviated ARIN meeting, typically held during a NANOG meeting.
PPML	Public Policy Mailing List, the ARIN email list dealing with policy issues.
PRDb	Policy-based Routing Database, the database used to configure inter-domain routing information for the NSFNET.
PSU	Public Sector Unit
PTCL	Pakistan Telecommunications Company, Ltd.
RA	Routing Arbiter, a NSF-funded organization acting as a neutral arbiter for inter-domain routing information on the Internet.
RADb	Routing Assets Database, the routing information database maintained by the Routing Arbiter.
RFC	Request For Comments, the document series started by the NWG and continued by the IETF to document their deliberations.
RIB	Routing Information Base
RIPE NCC	Réseaux Internet Protocol Européens Network Coordination Center, the European analogue to NANOG, also functioning as the European RIR.
RIR	Regional Internet Registry
RPKI	Resource Public Key Infrastructure, a family of IETF standards for securing BGP.
RWhois	Referral Whois
SANOG	South Asia Network Operators Group
SIDR WG	Secure Inter-Domain Routing Working Group, a working group at the IETF dedicated to creating a more secure version of BGP, which depends on RPKI.
STPI	Software Technology Parks of India
SWIP	Shared Whois Project
TCP	Transmission Control Protocol; originally Transmission Control Program.
TLD	Top-Level Domain
TRAI	Telecommunications Regulatory Authority of India
UDRP	Uniform Domain-Name Dispute Resolution Policy
UNDP	United Nations Development Programme

vBNS	Very high speed Backbone Network Services, a backbone network to support academic and research institutions after the NSFNET was privatized.
VoIP	Voice over IP
VSNL	Videsh Sanchar Nigam Ltd
WCIT	World Conference on International Telecommunications
WSIS	World Summit on the Information Society

Bibliography

- Abbate, J. (1999). *Inventing the Internet*. MIT Press.
- Abell, P. and Reyniers, D. (2000). On the Failure of Social Theory. *British Journal of Sociology*, 51(4):739–750.
- Agnew, J. (1994). The Territorial Trap: The Geographical Assumptions of International Relations Theory. *Review of International Political Economy*, 1(1):53–80. Available from: <http://www.tandfonline.com/doi/abs/10.1080/09692299408434268>.
- Allen, J. (2011). Topological Twists: Power’s Shifting Geographies. *Dialogues in Human Geography*, 1(3):283–298. Available from: <http://dhg.sagepub.com/lookup/doi/10.1177/2043820611421546>.
- Alvestrand, H. (2004). RFC 3935: A Mission Statement for the IETF. Available from: <http://tools.ietf.org/html/rfc3935>.
- Anderson, B. (2006). *Imagined Communities: Reflections on the Origin and Spread of Nationalism, New Edition*. Verso, 2nd edition.
- Arrighi, G. (1990). The Three Hegemonies of Historical Capitalism. *Review (Fernand Braudel Center)*, 13(3):365–408.
- Barlow, J. P. (1996). A Declaration of the Independence of Cyberspace. Available from: <http://homes.eff.org/~barlow/Declaration-Final.html>.
- Bell, D. (1973). *The Coming of Post-Industrial Society: A Venture in Social Forecasting*. Basic Books, New York, NY, USA.
- Benkler, Y. (2007). *The Wealth of Networks*. Yale University Press.
- Blanchette, J.-F. (2011). A Material History of Bits. *Journal of the American Society for Information Science and Technology*, 62(6):1042–1057.
- Boden, D. and Molotch, H. L. (1994). The Compulsion of Proximity. In Friedland, R. and Boden, D., editors, *NowHere: Space, Time and Modernity*, pages 257–286. University of California Press.

- Boothe, P., Hiebert, J., and Bush, R. (2006). How Prevalent is Prefix Hijacking on the Internet? In *Proceedings of NANOG36*. Available from: <http://www.nanog.org/meetings/abstract?id=411>.
- Bowker, G. C. (1994). Information Mythology and Infrastructure. In Bud-Frierman, L., editor, *Information Acumen: The Understanding and Use of Information in Modern Business*, pages 231–247. Routledge.
- Bowker, G. C. and Star, S. L. (2000). *Sorting Things Out: Classification and Its Consequences*. MIT Press.
- Brenner, N. (1999). Beyond State-Centrism? Space, Territoriality, and Geographical Scale in Globalization Studies. *Theory and Society*, 28(1):39–78. Available from: <http://www.jstor.org/stable/3108505>.
- Brown, J. S. and Duguid, P. (2000). *The Social Life of Information*. Harvard Business Press.
- Burrell, J. (2009). The Field Site as a Network: A Strategy for Locating Ethnographic Research. *Field Methods*, 21(2):181–199. Available from: <http://fm.sagepub.com/cgi/doi/10.1177/1525822X08329699>.
- Bush, R. (2005). Into the Future with the Internet Vendor Task Force, a Very Curmudgeonly View, or Testing Spaghetti: A Wall’s Point of View. *SIGCOMM Comput. Commun. Rev.*, 35(5):67–68. Available from: <http://portal.acm.org/citation.cfm?id=1096536.1096544>.
- Carey, J. W. and Quirk, J. J. (1988). The History of the Future. In *Communication as Culture: Essays on Media and Society*, pages 173–200. Unwin Hyman.
- Carr, S., Crocker, S., and Cerf, V. G. (1970). HOST-HOST Communication Protocol in the ARPA Network. In *Proc. AFIPS SJCC*, Los Alamitos, CA, USA. Available from: <http://doi.ieeecomputersociety.org/10.1109/AFIPS.1970.40>.
- Castells, M. (2000). *The Rise of the Network Society*. Wiley-Blackwell, 2nd edition.
- Cerf, V. G. (1990). Interview by Judy O’Neill.
- Cerf, V. G., Dalal, Y., and Sunshine, C. (1974). RFC 675: Specification of Internet Transmission Control Program. Available from: <http://tools.ietf.org/html/rfc675>.
- Cerf, V. G. and Mills, K. (1990). RFC 1169: Explaining the Role of GOSIP. Available from: <http://tools.ietf.org/search/rfc1169>.
- Chapin, A. L. (1992). RFC 1358: Charter of the Internet Architecture Board (IAB). Available from: <http://www.ietf.org/rfc/rfc1358.txt>.

- Chatzis, K. (1999). Designing and Operating Storm Water Drain Systems: Empirical Findings and Conceptual Developments. In Coutard, O., editor, *The Governance of Large Technical Systems*, pages 73–90. Routledge.
- Cheshire, C. (2011). Online Trust, Trustworthiness, or Assurance? *Daedalus*, 140(4):49–58.
- Chowdary, T. H. (2011). Birth and Growth of the Internet in India. In Rao, M. and Manzar, O., editors, *netCh@kra: 15 Years of the Internet in India <Retrospectives and Roadmaps>*, pages 63–72. Digital Empowerment Foundation.
- Christianson, B. and Harbison, W. S. (1997). Why Isn't Trust Transitive? In Lomas, M., editor, *Security Protocols: Lecture Notes in Computer Science*, volume 1189, pages 171–176. Springer Berlin/Heidelberg. Available from: <http://www.springerlink.com/index/3r64847m6w7256x1.pdf>.
- Clark, D. D. (1988). The design philosophy of the DARPA internet protocols. In *Symposium proceedings on Communications architectures and protocols*, pages 106–114, Stanford, California, United States. ACM. Available from: <http://portal.acm.org/citation.cfm?id=52336>http://portal.acm.org/ft_gateway.cfm?id=52336&type=pdf&coll=GUIDE&d1=GUIDE&CFID=68519331&CFTOKEN=59122825.
- Clark, D. D., Wroclawski, J., Sollins, K. R., and Braden, R. (2002). Tussle in Cyberspace: Defining Tomorrow's Internet. In *Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM '02)*, pages 347–356, New York, NY, USA. ACM. Available from: <http://portal.acm.org/citation.cfm?id=633025.633059>.
- Cohen, A. P. (1985). *The Symbolic Construction of Community*. Routledge.
- Cohen, J. (1999). Reflections on Habermas on Democracy. *Ratio Juris*, 12(4):385–416. Available from: <http://www.blackwell-synergy.com/links/doi/10.1111%2F1467-9337.00132>.
- Cook, K. S., Yamagishi, T., Cheshire, C., Cooper, R., Matsuda, M., and Mashima, R. (2005). Trust Building via Risk Taking: A Cross-Societal Experiment. *Social Psychology Quarterly*, 68(2):121–142. Available from: <http://spq.sagepub.com/cgi/doi/10.1177/019027250506800202>.
- Cotton, M. and Vegoda, L. (2010). RFC 5735/BCP 153: Special Use IPv4 Addresses. Available from: <http://tools.ietf.org/html/rfc5735>.
- Coutard, O., editor (1999). *The Governance of Large Technical Systems*. Routledge.
- Crocker, S. (1969). RFC 3: Documentation Conventions. Available from: <http://tools.ietf.org/html/rfc3>.

- Davies, E. and Hofmann, J. (2004). RFC 3844: IETF Problem Resolution Process. Available from: <http://tools.ietf.org/html/rfc3844>.
- DeNardis, L. and Raymond, M. (2013). Thinking Clearly about Multistakeholder Internet Governance. In *Proceedings of the 8th Annual GigaNet symposium*, Bali, Indonesia. Available from: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2354377.
- Donath, J. (1999). Identity and Deception in the Virtual Community. In Kollock, P. and Smith, M., editors, *Communities in Cyberspace*, pages 27–58. Routledge.
- Donley, C., Howard, L., Kuarsingh, V., Berg, J., and Doshi, J. (2013). RFC 7021: Assessing the Impact of Carrier-Grade NAT on Network Applications. Available from: <https://tools.ietf.org/html/rfc7021>.
- Drori, G. S. (2007). Information Society as a Global Policy Agenda: What Does It Tell Us About the Age of Globalization? *International Journal of Comparative Sociology*, 48(4):297–316. Available from: <http://cos.sagepub.com/cgi/content/abstract/48/4/297>.
- Edelman, B. (2009). Running Out of Numbers: Scarcity of IP Addresses and What To Do About It.
- Edelman, B. B. and Ryan, S. M. (2013). Guidance from ARIN on Legal Aspects of the Transfer of Internet Protocol Numbers. *Business Law Today*, (May):1–5. Available from: <http://apps.americanbar.org/buslaw/blt/content/2013/05/article-03-edelman.shtml>.
- Edwards, P. N. (2003). Infrastructure and Modernity: Force, Time, and Social Organization in the History of Sociotechnical Systems. In Misa, T. J., Brey, P., and Feenberg, A., editors, *Modernity and Technology*, pages 185–226.
- Edwards, P. N., Jackson, S. J., Bowker, G. C., and Knobel, C. P. (2007). Understanding Infrastructure: Dynamics, Tensions, and Design. Technical Report January, School of Information, University of Michigan, Ann Arbor.
- Egevang, K. and Francis, P. (1994). RFC 1631: The IP Network Address Translator (NAT). Available from: <http://www.ietf.org/rfc/rfc1631.txt>.
- Evans, P. B. (1995). *Embedded Autonomy: States and Industrial Transformation*. Princeton University Press.
- Farrell, H. (2009). Constructing Mid-Range Theories of Trust: The Role of Institutions. In Cook, K. S., Hardin, R., and Levi, M., editors, *Whom Can We Trust? How Groups, Networks, and Institutions Make Trust Possible*. Russell Sage Foundation, New York.

- Feamster, N., Balakrishnan, H., and Rexford, J. (2004). Some Foundational Problems in Interdomain Routing. In *3rd ACM SIGCOMM Workshop on Hot Topics in Networks (HotNets)*. Available from: <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.132.3571&rep=rep1&type=pdf>.
- Ferguson, P. and Senie, D. (2000). RFC2827/BCP38: Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing. Available from: <http://tools.ietf.org/html/bcp38>.
- Foster, J. B. and McChesney, R. W. (2011). The Internet's Unholy Marriage to Capitalism. *Monthly Review*, 62(10). Available from: <http://monthlyreview.org/2011/03/01/the-internets-unholy-marriage-to-capitalism/>.
- Galloway, A. R. (2005). Global Networks and the Effects on Culture. *Annals of the American Academy of Political and Social Science*, 597(1):19–31. Available from: <http://ann.sagepub.com/cgi/doi/10.1177/0002716204270066>.
- Gao, L. (2001). On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Transactions on Networking*, 9(6):733–745. Available from: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=974527>.
- Gao, L. and Rexford, J. (2001). Stable Internet Routing Without Global Coordination. *IEEE/ACM Transactions on Networking*. Available from: <http://dl.acm.org/citation.cfm?id=504612>.
- Garnham, N. (2000). 'Information Society' as Theory or Ideology. *Information, Communication and Society*, 3(2):139–152. Available from: <http://www.ingentaconnect.com/content/routledg/rics/2000/00000003/00000002/art00002>.
- Gellner, E. (1988). Trust, Cohesion, and the Social Order. In Gambetta, D., editor, *Trust: Making and Breaking Cooperative Relations*, pages 142–157. Basil Blackwell.
- Gerich, E. (1992). RFC 1366: Guidelines for Management of IP Address Space. Available from: <http://tools.ietf.org/html/rfc1366>.
- Giddens, A. (1991). *The Consequences of Modernity*. Stanford University Press.
- Gill, P., Arlitt, M., Li, Z., and Mahanti, A. (2008). The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse? In *Proceedings of Passive and Active Measurement (PAM) Conference 2008*, Cleveland, USA.
- Goldsmith, J. and Wu, T. (2006). *Who Controls the Internet?: Illusions of a Borderless World*. Oxford University Press, USA.
- Granovetter, M. (1985). Economic Action and Social Structure: The Problem of Embeddedness. *American Journal of Sociology*, 91(3):481–510. Available from: <http://www.jstor.org/stable/2780199>.

- Haas, P. M. (1992). Introduction: Epistemic Communities and International Policy Coordination. *International Organization*, 46(1):1–35.
- Hafner, K. and Lyon, M. (1998). *Where Wizards Stay Up Late: The Origins Of The Internet*. Simon & Schuster. Available from: <http://www.amazon.com/dp/0684832674>.
- Hardin, R. (2002). *Trust and Trustworthiness*. Russell Sage Foundation Publications.
- Harris, S. R. and Gerich, E. (1996). Retiring the NSFNET Backbone Service: Chronicling the End of an Era. *Connexions*, 10(4). Available from: http://www.merit.edu/networkresearch/projecthistory/nsfnet/nsfnet_article.php.
- Harvey, D. (1985). The Geopolitics of Capitalism. In Gregory, D. and Urry, J., editors, *Social Relations and Spatial Structures*, pages 128–163. Palgrave Macmillan.
- Harvey, D. (1990). *The Condition of Postmodernity: An Enquiry into the Origins of Cultural Change*. Wiley-Blackwell.
- Harvey, D. (2007). *Limits to Capital*. Verso, 2nd edition.
- Heiskala, R. (2003). Informational Revolution, the Net and Cultural Identity: A Conceptual Critique of Manuel Castells’s The Information Age. *European Journal of Cultural Studies*, 6(2):233–245. Available from: <http://ecs.sagepub.com>.
- Hoffman, P. (2012). The Tao of IETF: A Novice’s Guide to the Internet Engineering Task Force. Available from: <http://www.ietf.org/tao.html>.
- Housely, R., Curran, J., Huston, G., and Conrad, D. (2013). RFC 7020: The Internet Numbers Registry System. Available from: <http://tools.ietf.org/html/rfc7020>.
- Hovey, R. and Bradner, S. (1996). RFC 2028: The Organizations Involved in the IETF Standards Process. Available from: <http://tools.ietf.org/html/bcp11>.
- Hubbard, K., Kusters, M., Conrad, D., Karrenberg, D., and Postel, J. (1996). RFC 2050/BCP 12: Internet Registry IP Allocation Guidelines. Available from: <http://tools.ietf.org/html/rfc2050>.
- Hughes, T. P. (1983). *Networks of Power: Electrification in Western Society, 1880-1930*. Johns Hopkins University Press.
- Huizer, E. (1996). RFC 2031: IETF-ISOC relationship. Available from: <http://tools.ietf.org/html/rfc2031>.
- IAB Advisory Committee (2004). RFC 3716: The IETF in the Large: Administration and Execution. Available from: <http://tools.ietf.org/html/rfc3716>.
- Innis, H. (1951). *The Bias of Communication*. University of Toronto Press.

- Internet Society (2011). Internet Society Background Paper: International Telecommunications Regulations. Available from: <http://internetsociety.org/background-international-telecommunication-regulations>.
- Jackson, S. J., Edwards, P. N., Bowker, G. C., and Knobel, C. P. (2007). Understanding Infrastructure: History, Heuristics and Cyberinfrastructure Policy. *First Monday*, 6(4). Available from: <http://firstmonday.org/ojs/index.php/fm/article/view/1904/1786>.
- Jessop, B. (2007). Knowledge as a Fictitious Commodity: Insights and Limits of a Polanyian Perspective. In Bugra, A. and Agartan, K., editors, *Reading Karl Polanyi for the Twenty-First Century: Market Economy as Political Project*, pages 115–134. Palgrave Macmillan.
- Jessop, B., Brenner, N., and Jones, M. (2008). Theorizing Sociospatial Relations. *Environment and Planning D: Society and Space*, 26(3):389 – 401. Available from: <http://www.envplan.com/abstract.cgi?id=d9107>.
- Joerges, B. (1999). Do Politics Have Artifacts? *Social Studies of Science*, 29(3):411–431.
- Johnson, D. R., Crawford, S. P., and Palfrey, J. G. (2004). The Accountable Net: Peer Production of Internet Governance. *Virginia Journal of Law and Technology*, 9(9). Available from: <http://ssrn.com/abstract=529022>.
- Khare, V., Ju, Q., and Zhang, B. (2012). Concurrent Prefix Hijacks: Occurrence and Impacts. In *Proceedings of the 2012 ACM Conference on Internet Measurement*, pages 29–35. Available from: <http://dl.acm.org/citation.cfm?id=2398780>.
- Kollock, P. (1994). The Emergence of Exchange Structures: An Experimental Study of Uncertainty, Commitment, and Trust. *American Journal of Sociology*, 100(2):313–345. Available from: <http://www.jstor.org/stable/2782072>.
- Krol, E. and Hoffman, E. (1993). RFC 1462: FYI on "What is the Internet?". Available from: <http://tools.ietf.org/html/rfc1462.html>.
- Kuerbis, B. (2011). *Securing Critical Internet Resources: Influencing Internet Governance through Social Networks and Delegation*. PhD thesis, Syracuse University. Available from: http://surface.syr.edu/it_etd/68/.
- Labovitz, C., Lekel-Johnson, S., McPherson, D., Oberheide, J., and Jahanian, F. (2010). Internet Inter-Domain Traffic. In *Proceedings of the ACM SIGCOMM 2010 Conference*, pages 75–86, New Delhi, India. Available from: http://ccr.sigcomm.org/online/?q=node/667http://ccr.sigcomm.org/online/files/p75_0.pdf.
- Lave, J. and Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press.

- Lefebvre, H. (1991). *The Production of Space*. Wiley-Blackwell.
- Lefebvre, H. (2003). Space and the State. In Brenner, N., Jessop, B., Jones, M., and Macleod, G., editors, *State/Space: A Reader*, pages 84–100. Wiley-Blackwell.
- Lewis, J. D. and Weigert, A. (1985). Trust as a Social Reality. *Social Forces*, 63(4):967–985. Available from: <http://www.jstor.org/stable/2578601>.
- Li, L., Alderson, D., Willinger, W., and Doyle, J. (2004). A First-principles Approach to Understanding the Internet's Router-level Topology. *ACM SIGCOMM Computer Communication Review*, 34(4):3–14. Available from: <http://portal.acm.org/citation.cfm?doid=1030194.1015470>.
- Lougheed, K. and Rekhter, Y. (1989). RFC 1105: A Border Gateway Protocol. Available from: <http://www.ietf.org/rfc/rfc1105.txt>.
- Luhmann, N. (1979). *Trust and Power*. John Wiley and Sons.
- Luhmann, N. (1988). Familiarity, Confidence, Trust: Problems and Alternatives. In Gambetta, D., editor, *Trust: Making and Breaking Cooperative Relations*, pages 94–107. Basil Blackwell.
- Lyotard, J.-F. (1985). *The Post-Modern Condition*. University of Minnesota Press, Minneapolis.
- Mann, M. (1984). The Autonomous Power of the State: Its Origins, Mechanisms and Results. *European Journal of Sociology*, 25(2):185–213.
- Marcus, G. E. (1995). Ethnography in/of the World System: The Emergence of Multi-Sited Ethnography. *Annual Review of Anthropology*, 24:95–117. Available from: <http://arjournals.annualreviews.org/doi/abs/10.1146/annurev.an.24.100195.000523>.
- Marx, L. (1964). *The Machine in the Garden: Technology and the Pastoral Ideal in America*. Oxford University Press.
- Massey, D. (1994a). A Global Sense of Place. In *Space, Place and Gender*, pages 146–156. University of Minnesota Press.
- Massey, D. (1994b). *Space, Place, and Gender*. University of Minnesota Press.
- Massey, D. (1994c). Uneven Development: Social Change and Spatial Divisions of Labour. In *Space, Place and Gender*, pages 86–114. University of Minnesota Press.
- Mathew, A. J. and Cheshire, C. (2010). The New Cartographers: Trust and Social Order within the Internet Infrastructure. In *Proceedings of the 38th Research Conference on Communication, Information and Internet Policy (Telecommunications and Policy Research Conference)*, George Mason University School of Law, Arlington, VA. Available from: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1988216.

- Mattelart, A. (2003). *The Information Society: An Introduction*. Sage Publications.
- Mayntz, R. and Hughes, T. P., editors (1988). *The Development of Large Technical Systems*. Campus Verlag; Westview Press.
- McPherson, D. and Patel, K. (2006). RFC 4277: Experience with the BGP-4 Protocol. Available from: <http://tools.ietf.org/html/rfc4277>.
- Mueller, M. L. (2002). *Ruling the Root: Internet Governance and the Taming of Cyberspace*. MIT Press.
- Mueller, M. L. (2010). *Networks and States: The Global Politics of Internet Governance*. MIT Press.
- Mueller, M. L., Schmidt, A., and Kuerbis, B. (2013). Internet Security and Networked Governance in International Relations. *International Studies Review*, 15:86–104. Available from: <http://doi.wiley.com/10.1111/misr.12024>.
- Nissenbaum, H. (2004). Will Security Enhance Trust Online, or Supplant It? In Roderick, K. M. and Cook, K. S., editors, *Trust and Distrust in Organizations: Dilemmas and Approaches*, pages 155–188. Russell Sage Foundation Publications. Available from: <http://www.nyu.edu/projects/nissenbaum/papers/trust.pdf>.
- NSF (1993). NSF 93-52 - Network Access Point Manager, Routing Arbiter, Regional Network Providers, and Very High Speed Backbone Network Services Provider for NSFNET and the NREN (SM) Program. Available from: http://w2.eff.org/Infrastructure/NREN_NSFNET_NPN/nsf_nren.rfp.
- Oliveira, R. and Willinger, W. (2010). The (In)Completeness of the Observed Internet AS-level Structure. *IEEE/ACM Transactions on Networking*, 18(1):109–122. Available from: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5200324>.
- Oliveira, R. V., Pei, D., Willinger, W., Zhang, B., and Zhang, L. (2008). In search of the elusive ground truth: the Internet’s AS-level connectivity structure. *SIGMETRICS Perform. Eval. Rev.*, 36(1):217–228. Available from: <http://portal.acm.org/citation.cfm?id=1375482>.
- Piscitello, D. M. and Chapin, A. L. (1993). *Open Systems Networking: TCP/IP and OSI*. Addison-Wesley.
- Polanyi, K. (2001). *The Great Transformation*. Beacon Press, 2nd edition.
- Polanyi, M. (1966). *The Tacit Dimension*. Doubleday.
- Postel, J. (1981). RFC 791: Internet Protocol. Available from: <http://tools.ietf.org/html/rfc791>.

- Powell, W. W. (1990). Neither Market nor Hierarchy: Network Forms of Organization. *Research in Organizational Behavior*, 12:295–336.
- Ramachandran, A. and Feamster, N. (2006). Understanding the Network-level Behavior of Spammers. *ACM SIGCOMM Computer Communication Review*, 36(4):291. Available from: <http://portal.acm.org/citation.cfm?doid=1151659.1159947>.
- Ramani, S. (2011). Bringing the Internet to India. In Rao, M. and Manzar, O., editors, *netCh@kra: 15 Years of the Internet in India <Retrospectives and Roadmaps>*, pages 47–62. Digital Empowerment Foundation.
- Rekhter, Y. and Li, T. (1995). RFC 1771: A Border Gateway Protocol 4 (BGP-4). Available from: <http://tools.ietf.org/html/rfc1771>.
- Rekhter, Y., Li, T., and Hares, S. (2006). RFC 4271: A Border Gateway Protocol 4. Available from: <http://tools.ietf.org/html/rfc4271>.
- Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J., and Lear, E. (1996). RFC 1918/BCP 5: Address Allocation for Private Internets. Available from: <http://tools.ietf.org/html/rfc1918>.
- Rheingold, H. (1995). Introduction. The MIT Press. Available from: <http://www.rheingold.com/vc/book/intro.html>.
- Rose, G. (2011). The Economics of Internet Interconnection: Insights from the Comcast-Level3 Peering Dispute. Technical report, Federal Communications Commission. Available from: <http://apps.fcc.gov/ecfs//document/view.action?id=7021239457>.
- Rosen, E. C. (1982). RFC 827: Exterior Gateway Protocol. Available from: <http://tools.ietf.org/html/rfc827>.
- Russell, A. (2006). 'Rough Consensus and Running Code' and the Internet-OSI Standards War. *IEEE Annals of the History of Computing*, 28(3):48–61. Available from: <http://ieeexplore.ieee.org/stampPDF/getPDF.jsp?tp=&arnumber=1677461&isnumber=35276>.
- Sandvig, C. (2013). The Internet as Infrastructure. In *The Oxford Handbook of Internet Studies*, volume 1, pages 86–108. Oxford University Press.
- Schuchard, M., Mohaisen, A., Kune, D. F., Hopper, N., and Vasserman, E. Y. (2010). Losing Control of the Internet : Using the Data Plane to Attack the Control Plane. In *Proceedings of the 17th ACM conference on Computer and Communications Security (CCS '10)*, pages 726–728, New York, NY, USA. ACM. Available from: <http://doi.acm.org/10.1145/1866307.1866411>.

- Sheppard, E. (2002). The Spaces and Times of Globalization: Place, Scale, Networks, and Positionality. *Economic Geography*, 78(3):307–330. Available from: <http://www.jstor.org/stable/4140812>.
- Singhal, A. (2011). India's Internet Policy and Regulatory Regime. In Rao, M. and Manzar, O., editors, *netCh@kra: 15 Years of the Internet in India <Retrospectives and Roadmaps>*, pages 73–84. Digital Empowerment Foundation.
- Soja, E. W. (1989). *Postmodern Geographies: The Reassertion of Space in Critical Social Theory*. Verso.
- Star, S. L. (1999). The Ethnography of Infrastructure. *American Behavioral Scientist*, 43(3):377–391. Available from: <http://abs.sagepub.com/cgi/content/abstract/43/3/377>.
- Star, S. L. and Ruhleder, K. (1996). Steps Toward an Ecology of Infrastructure: Design and Access for Large Information Spaces. *Information Systems Research*, 7(1):111–134.
- Stoll, C. (2005). *The Cuckoo's Egg: Tracking a Spy through the Maze of Computer Espionage*. Gallery Books.
- Telecom Regulatory Authority of India (2014). The Indian Telecom Services Performance Indicators, October - December, 2013. Technical report, Telecom Regulatory Authority of India, New Delhi, India. Available from: http://www.trai.gov.in/Content/PerformanceIndicatorsReports/1_1_PerformanceIndicatorsReports.aspx.
- Thrift, N. (2004). Remembering the Technological Unconscious by Foregrounding Knowledges of Position. *Environment and Planning D: Society and Space*, 22(1):175–190.
- Turner, F. (2006a). *From Counterculture to Cyberculture: Stewart Brand, the Whole Earth Network, and the Rise of Digital Utopianism*. University Of Chicago Press.
- Turner, F. (2006b). How Digital Technology Found Utopian Ideology: Lessons From the First Hackers' Conference. In *Critical Cyberculture Studies*, chapter 22, pages 257–269. New York University Press. Available from: <http://www.stanford.edu/~fturner/TurnerHackersConferenceChapter.pdf>.
- United States Government (2009). Cyberspace Policy Review: Assuring a Trusted and Resilient Information and Communications Infrastructure. Technical report. Available from: http://www.whitehouse.gov/assets/documents/Cyberspace_Policy_Review_final.pdf.
- Villamizar, C., Chandra, R., and Govindan, R. (1998). RFC 2439: BGP Route Flap Damp-ing. Available from: <http://tools.ietf.org/html/rfc2439>.

- Weil, J., Kuarsingh, V., Donley, C., Liljenstolpe, C., and Azinger, M. (2012). RFC 6598: IANA-Reserved IPv4 Prefix for Shared Address Space. Available from: <http://tools.ietf.org/html/rfc6598>.
- Weller, D. and Woodcock, B. (2013). Internet Traffic Exchange: Market Developments and Policy Challenges. Available from: http://www.oecd-ilibrary.org/science-and-technology/internet-traffic-exchange_5k918gpt130q-en.
- White, R., McPherson, D., and Sangli, S. (2004). *Practical BGP*. Addison-Wesley Professional.
- Wilson, S. M. and Peterson, L. C. (2002). The Anthropology of Online Communities. *Annual Review of Anthropology*, 31(1):449–467. Available from: <http://www.annualreviews.org/doi/abs/10.1146/annurev.anthro.31.040402.085436>.
- Winner, L. (1980). Do Artifacts Have Politics? *Daedalus*, 109(1):121–136. Available from: <http://www.jstor.org/stable/20024652>.
- Yamagishi, T. and Yamagishi, M. (1994). Trust and Commitment in the United States and Japan. *Motivation and Emotion*, 18(2):129–166.
- Zhang, Y., Mao, Z., and Wang, J. (2007). Low-Rate TCP-Targeted DoS Attack Disrupts Internet Routing. In *In Proceedings of the 14th Annual Network & Distributed System Security Symposium*. Available from: <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.137.5004&rep=rep1&type=pdf>.