Pass the Packet, Please?

Explorations into Trust and the Social Organisation of Network Administration

Master's Thesis

Ashwin J. Mathew
School of Information, UC Berkeley
ashwin@ischool.berkeley.edu

Table of Contents

List of Abbreviations and Acronyms

ANS	Advanced Network and Services
ARIN	American Registry for Internet Numbers
ARPA	Advanced Research Projects Agency of the United State's Department of Defence
ARPANET	Research network funded by ARPA
AS	Autonomous System
ASN	Autonomous System Number
BARRNET	Bay Area Regional Research Network
BGP	Border Gateway Protocol
IANA	Internet Assigned Numbers Authority
IETF	Internet Engineering Task Force
IP	Internet Protocol
IRR	Internet Routing Registry
ISP	Internet Service Provider
NANOG	North American Network Operators' Group
NSFNET	Research network funded by the National Science Foundation of the USA
PIE	Pakistan Internet Exchange
RA	Routing Arbiter
RADb	Routing Assets Database
RIPE	Réseaux IP Européens, the European RIR
RIR	Regional Internet Registry
TCP	Transmission Control Protocol

Pass the Packet, Please?

Explorations into Trust and the Social Organisation of Network Administration[1]

1 Introduction

The most significant technologies in our lives are also the ones we take most for granted. Roads, water, gas: all these fade into the background, only to be noticed when they fail. In a word, these systems have become infrastructures (Edwards 2003; Star 1999). The Internet is the most recent of these, and is of particular interest for the speed with which it has gone from novelty to an aspect of everyday life. It is also, perhaps, the most invisible of infrastructures; we can see roads, we can understand where a water main breaks, but it is beyond our grasp to conceptualise flows of data on the Internet.

Once we view the Internet as built environment, rather than an abstract “cloud”, it becomes critical to understand the politics of its technological infrastructure, to engage with the particulars of its systems. There is a significant body of work devoted to the economic and political aspects of this infrastructure. In this work, I follow a different line of inquiry: the study of the social aspects of the Internet's infrastructure.

My focus here is on the social organisation of network administrators who maintain the interconnections between the many networks that make up the Internet. I analyse their

[1] This work would not have been possible without the generosity and expertise of the technologists I interviewed in the course of my research. I hope I have done justice to the rich and detailed conversations I was able to have with each of them. I am also grateful to Professor Coye Cheshire for his support and advice throughout this research. And finally, to Anasuya Sengupta, who kept me sane through this process, and edited it all when it was done.

interactions with one another in relation to the Border Gateway Protocol (BGP), a key Internet technology that mediates the process of network interconnection. Although BGP was created for the political purpose of allowing different network domains (typically separate commercial entities) to coexist, it paradoxically embeds notions of trust. This makes it a fascinating subject for study, to understand how trust came to be embedded in the technology, and to engage with the question of how this mediates the social organisation of network administration.

The arguments I present here are twofold. First, that social and cultural norms are embedded in the technologies we build, and that this embedding has material effects. Second, that the Internet embeds norms of trust, and that this has significant implications for the ways that we might think about the functioning of this artifact: in patterns of spread, in institutional arrangements, and in technological practices and design. I follow a conception of trust as being a relational quantity, established collectively, rather than individually. As Lewis and Weigert (1985, 968) put it:

> From a sociological perspective, trust must be conceived as a property of collective units (ongoing dyads, groups, and collectivities), not of isolated individuals. Being a collective attribute, trust is applicable to the relations among people rather than to their psychological states taken individually.

The remainder of this paper could be read as a gradual unfolding from the technological to the social. In Section 2, I describe the form and function of the Border Gateway Protocol, and raise the key research questions arising from the particular form that it has taken. In Section

3, I engage with the methodological and conceptual concerns that arose from this study. Section 4 deals with the evolution of the technology and the social groups, both of which are the subjects of this research. A more detailed examination of the ways in which trust and reputation manifest in this system is in Section 5. Finally, Section 6 concludes the paper, offering some final thoughts and possible directions for future research.

2 The (Infra)Structure of the Internet

The Internet is an instance of an "internet", a network of networks. Each of these networks optimizes data flows within itself across various technical parameters, such as overall bandwidth utilisation, or latency across different links in the network. The Internet was designed with this purpose in mind, to allow individual network administrations to operate as they pleased internally, but still interconnect to allow data to flow from one network to another, stitching these disparate networks into a seamless whole.

The structure of networks that forms as a result is essentially commercial: companies operating networks negotiate the terms of interconnection with one another. 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.

Inter-domain routing is the term used to refer to the process of routing data between different networks, managed as distinct administrative domains. Piscitello and Chapin 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.

As this suggests, it is practically impossible for every network to link directly to every other network on the Internet. Individual networks may operate in different geographies, and at different scales; some networks may offer local coverage within a city, or a region, while others may provide long distance links, across or between continents. Traffic from one network may have to transit multiple intermediate networks before reaching its destination.

Depending on their reach, networks are loosely classified as Tier 1, Tier 2, or stub networks. There are between 8 and 20 Tier 1 networks, depending on the specific definition of Tier 1. It is generally agreed that Tier 1 networks maintain a full routing table: a list of routes to all destinations on the Internet, currently numbering close to 300,000 (Huston 2009). In addition, Tier 1 networks maintain settlement-free peering agreements to all their peers, i.e., they carry sufficient volumes of traffic to one another that they agree not to pay one another for the privilege of interconnection. The contention about the exact number of Tier 1 networks is related to this issue, since some of these networks pay for interconnection to one or more of their peer networks. Tier 1 networks typically have global reach, and include companies such as AT&T, Sprint and Verizon. Tier 2 networks offer service to particular geographic regions, paying one or more Tier 1 networks for transit, and sometimes interconnecting with one another directly to avoid paying a Tier 1 for transit to a neighbouring Tier 2. Stub networks are those that exist at the periphery of the Internet, buying transit from Tier 1 or Tier 2 networks.

As such, the Internet is a complex graph of relationships amongst networks, loosely structured as a mesh of Tier 1 networks in the core, with Tier 2 networks arranged around these, feeding out to stub networks at the periphery (Illustration 1).

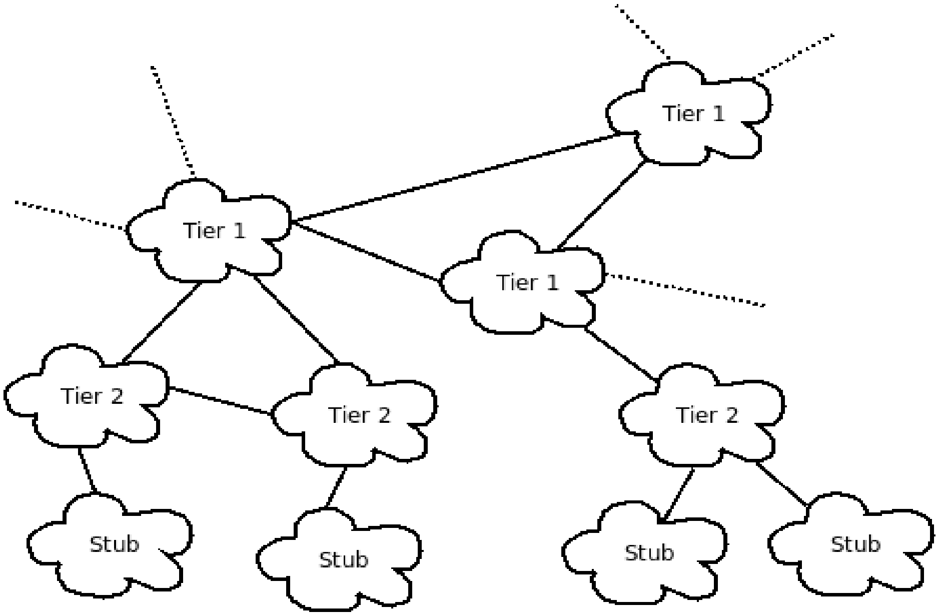

Illustration 1: Schematic of Tier 1, Tier 2 and Stub network interconnections

A useful analogy could be drawn to the postal system[2]. The address that a packet is destined for is stamped on the packet, but the specific route that the packet takes to its destination is dependent on the knowledge that different post offices have of one another. For instance, consider a packet destined for the city of Bangalore in the state of Karnataka in India, sent from the city of Berkeley in the state of California in the United States of America. Each relaying post office makes a local decision about where the packet should be sent next: the Berkeley post office may not know where Bangalore or Karnataka are, but knows that packets for India must be sent to the San Francisco post office. The San Francisco post office may similarly make a decision to send the packet to San Francisco International Airport, where a decision may be taken to forward the packet to Mumbai International Airport in India, and so on, until the packet reaches its destination. This is a fairly straightforward description of the

[2] This is not to say that this is how the postal system actually functions; I offer this as a useful image for clarity.

packet-switching technology of the Internet: packets are stamped with the Internet Protocol (IP) address of their destination, and routers (the post office analogue) make local decisions about where the packet should be sent.

Two important elements of complexity that are present in the Internet are missing from this analogy. First, that routers – the hardware responsible for moving data through a network – may be aware of multiple possible routes to the same destination, due to redundant paths to the destination through multiple neighbour (or “peer”) networks. Second, the mechanism used by networks to discover the destinations to which their peers are capable of carrying packets.

My research is particularly concerned with the technology used to address the latter source of complexity, the Border Gateway Protocol (BGP), which is used by networks to “advertise” to their peers the destinations to which they can carry routes (Rekhter and Li 1995). Although not as well known as its cousins, TCP/IP[3], BGP is one of the Internet's key infrastructural technologies, playing an essential role in inter-domain routing.

BGP uses the term “autonomous system” or “AS” to refer to a network. Some larger network operators may, in fact, manage multiple autonomous systems as part of their network infrastructure. Each autonomous system is assigned a unique autonomous system number (ASN) by a Regional Internet Registry (RIR)[4], which in turn is granted blocks of ASNs for allocation by the Internet Assigned Numbers Authority (IANA). For instance, the autonomous

[3] Transmission Control Protocol/Internet Protocol, two of the foundational elements of the Internet's technology, and frequently mentioned in the popular press.

[4] There are currently 5 RIRs for different geographic regions: RIPE for Europe, AfriNIC for Africa, APNIC for Asia-Pacific, LACNIC for Latin America and the Caribbean and ARIN for North America.

system operated by the University of California, Berkeley is AS25 (Autonomous System number 25). There are currently a little over 31,000 autonomous systems active on the Internet (Huston 2009). The complexity of the graph of autonomous systems is illustrated in this visualisation (Illustration 2) from the Cooperative Association for Internet Data Analysis (CAIDA), which shows the interconnections between autonomous systems, and provides a sense of the hierarchical form of the Internet, with a cluster of Tier 1 networks in the core, leading out to Tier 2 networks, and then on to stub networks.

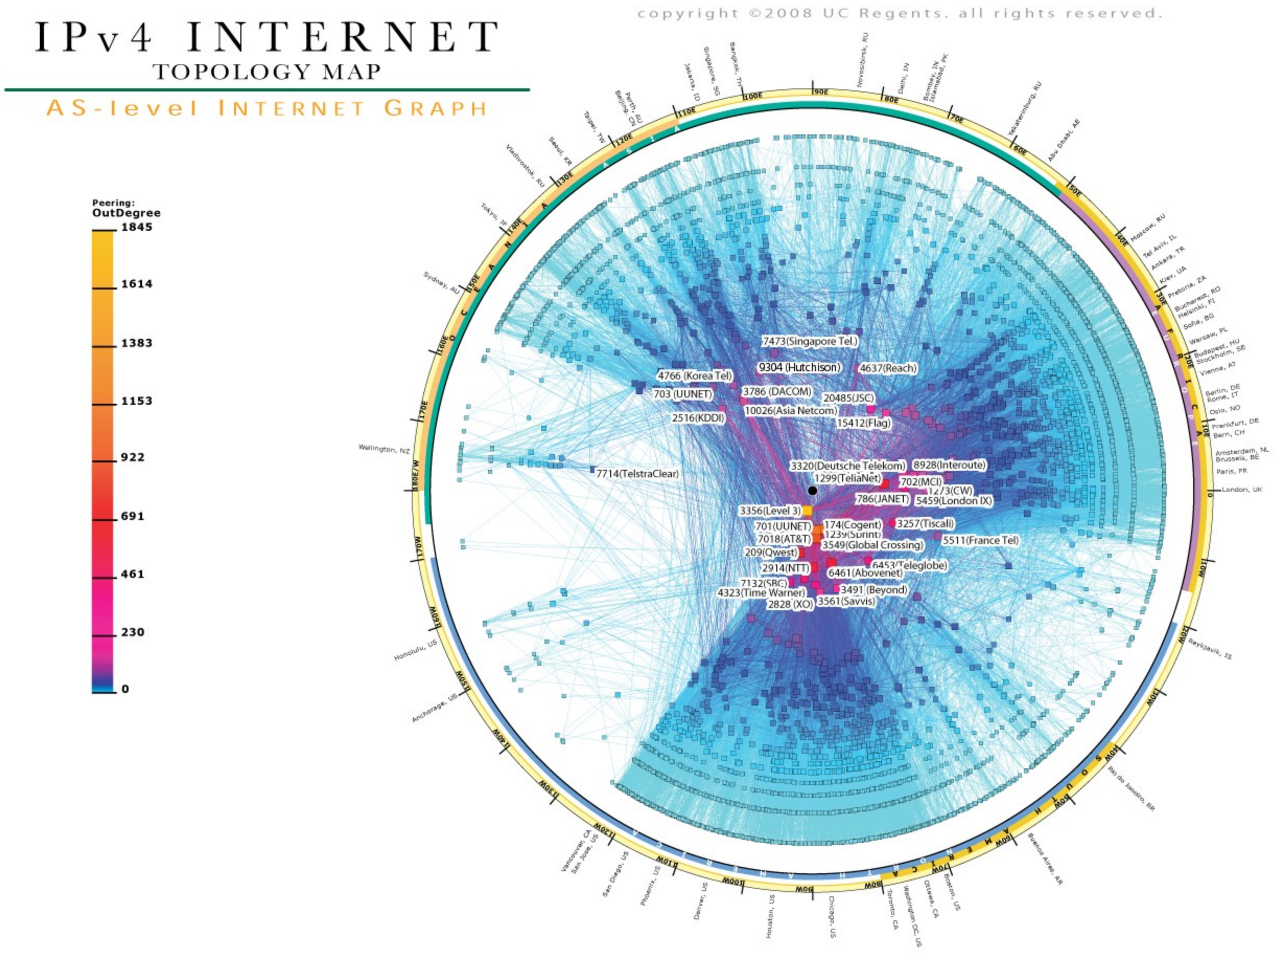

Illustration 2: CAIDA AS Level Internet Graph (Huffaker and claffy 2008)

The RIRs are also responsible for assigning blocks of IP address space, specified by IP address prefixes, to ASNs. This situation is made somewhat more complex by the fact that customer networks may lease the rights to an IP address prefix from their transit provider network.

A prefix is analogous to a locality from the post office example, but it is perhaps better illustrated using the example of a phone number. If we were to break down the phone number +91-80-2304537, “+91” is a prefix that specifies the country India, “+91-80” is a prefix that specifies the city of Bangalore in India, and so on. The longer the prefix, the more specific it is. Also, since IP addresses are of a fixed length (4 bytes), a longer prefix implies fewer unique IP addresses that are available behind that prefix.

BGP allows autonomous systems to “advertise” the prefixes that they can route to their peers. Their peers, in turn, forward these advertisements on to *their* peers, and so on, propagating information about which autonomous systems contain which IP address prefixes. In addition, this propagation of advertisements also functions as routing information: an AS receiving an advertisement for a prefix from one of its peers treats that advertisement as a claim that that peer knows how to get to a particular prefix. It will treat that peer as a possible route for data destined for an IP address contained within the advertised prefix. If a network is has multiple peers, it is entirely possible that it will receive advertisements for the same prefix from different peers, offering different paths to the same destination. Under such circumstances, the border routers in a network are configured to prefer one path or another depending on different parameters, such as the cost for transit across each of these paths.

The description that I've provided so far has been largely technological, with seemingly little of sociological interest. However, BGP paradoxically embeds notions of trust, which is surprising for a protocol that was created for the political purpose of allowing separate network administrations – *autonomous* systems – to coexist and interconnect with one another. That is to say, as a technology, it provides no means for autonomous systems to assess the credibility of advertisements received from their neighbours.

When an AS receives a route advertisement from one of its peers, it must trust that that peer is telling the truth. There is no mechanism for an AS to determine whether or not the claim that the peer makes to be able to carry traffic to a particular destination is valid or not. Indeed, there is not even a mechanism for validating that the autonomous system which originated an advertisement (remember, advertisements are often forwarded across multiple autonomous systems) actually has the rights to advertise that IP address space. This is in part a historical issue, since the records of assignment of prefixes to autonomous systems are themselves not entirely trustworthy, an issue which I will return to in Section 4. More importantly, this is a problem with BGP itself, as the protocol offers no support for validating the claims that peers make in their advertisements. This problem is magnified many times over by the fact that peers often act simply as relays for advertisements from *their* peers. It is, more often than not, the case that the AS originating an advertisement will have no commercial peering relationship whatsoever with the AS that receives it, since the advertisement may have transited multiple autonomous systems en route.

As a result of the trust embedded in BGP, it is possible for networks to make false claims about the routes that they are capable of carrying, either through inadvertent misconfiguration, or through malice. Google's YouTube was a recent high profile victim of such a false advertisement. The Government of Pakistan issued an order banning a particular video on YouTube; the Pakistan Internet Exchange (PIE) implemented this order through a BGP configuration, issuing an advertisement claiming that they knew how to get to YouTube, and then dropping all traffic that was directed along this route. Unfortunately, this advertisement propagated outside Pakistan, through one of PIE's peers in Singapore, and then on to the rest of the Internet. It took about 2 hours for the situation to be resolved, during which time most of the Internet believed YouTube to be located in Pakistan. Furthermore, PIE was disconnected from its peers during this period, effectively disconnecting Pakistan from the Internet (Brown 2008).

This is by no means an isolated incident. Similar, and sometimes even worse, incidents have occurred over the years. In addition, “prefix hijack” is an ongoing problem, where attackers maliciously send out false advertisements to draw traffic meant for a particular network towards themselves.

The principal tool that networks have at their disposal to limit these problems is the route filter. These filters are applied both when a network advertises a prefix (to filter which prefixes are advertised to which peers) and when a network receives an advertisement (to filter the prefixes that the network will route for a peer). Route filters are often applied for commercial purposes: for instance, a network with multiple peers may not want to act as a conduit for traffic between its peers, and so may configure filters to ensure that it accepts its peers' advertisements, but does not relay them. Tier 1 and Tier 2 networks may use filters to limit customers to advertising only the IP address prefixes which they have rights to, mitigating the problem of false advertisements to some extent. However, this only works well with immediate peers; when a peer is just acting as a relay for an advertisement, route filtering is not as effective, given the complex graph of relationships across which an advertisement may flow. In addition, interconnections amongst Tier 1 networks, and many Tier 2 networks, carry such volumes of traffic, to such a wide range of addresses, that it is simply impractical to filter route advertisements.

The trust manifest in BGP is not, therefore, simply the trust expressed between peers, but is *transitive trust* that establishes itself across the web of interconnected networks that make up the Internet. To trust a peer network is to trust that network's peers, and so on. This leads to the two principal research questions that I engage with in this work:

- How did BGP come to embed notions of trust?
- Given that BGP is trusting, how do network administrators communicate and coordinate to maintain the inter-domain routing system?

3 Constructing the Network Field

I began this research with an implicit assumption, that the use of the term “network” would refer to networking in the computer science sense of the word, in reference to technological infrastructure. I had thought to disambiguate this from networks of people by using the term “social network”. As may be evident from my research questions, I was in search of the social organisation of network administrators, as separate from the technical and commercial structure of the Internet.

While a useful distinction to make for directing this work, I found myself using “network” in multiple senses, as my research progressed: as a technical artifact managed by one or more administrators, as the collection of interconnected networks comprising the Internet, and, in less obvious ways, as an interconnected system of humans and objects. It is difficult to disambiguate these uses, as many of the social ties that I examine are explicitly mediated by the process of management of interconnections between two networks (in the Internet sense), and also expressed as a concern for the well-being of the Internet as a whole.

Conceptually, therefore, network administrators exist in a recursive relationship with the networks they manage, and the Internet as a whole. Networks cannot exist without network administrators. Equally, the role of a network administrator is defined by the existence of a network to be administered. To be a network administrator, particularly one managing BGP routers, is to enter into relationships with a network, or networks, with the concept of the Internet as a whole, and with other administrators, both in peer networks, and also as part of a larger community.

An additional important consideration arises from the fact that network administrators who manage BGP routers are doing so at the border between their network and their peers' networks. These administrators are physically remote from one another, bound by technical and commercial arrangements. Administrators also have a sense of being part of a larger community, as a consequence of professional practice, due to the transitive nature of trust embedded in BGP, and other security and coordination issues in the Internet beyond their own networks. This larger community of network administrators connects via mailing lists and occasional conferences, and also through a shared sense of history, passed on both textually and orally. As a result, there is no single fieldsite, or bounded set of fieldsites, that might be said to constitute the field. In essence, the network *is* the field.

I take these two features – recursive relationships amongst human and technological elements, and the lack of a bounded set of fieldsites – to be the defining characteristics of this kind of field, which I term a *network field*[5]. Each of these features raise particular conceptual, methodological and epistemological challenges, which also present the opportunity for critiquing and extending conventional research methods and modes of analysis.

3.1 Interpreting Socio-technical Systems

The problem of conceptualising socio-technical systems has been a recurring theme in the field of Science and Technology studies, with a general recognition that that the social and the technological are mutually constituted, rather than distinct elements to be studied and theorised independent of one another (Suchman 2009).

[5] In the absence of a term that encompasses the complexities I am describing, I have created the term “network field” as a useful heuristic to describe my research field.

Actor-Network Theory approaches this issue by positing that society is a network of heterogeneous materials comprising links amongst human and non-human elements, in which neither kind of element should be privileged in the investigation and analysis of a setting (Law 1992). Latour (2005) cautions us not to go in search of social groups, but rather to look into group formation, asking questions of the dynamics within groups, examining the relative coherence of a group as a process, rather than a static configuration. These ideas suggest the stance of investigating the Internet's inter-domain routing system not just as constituted of networks and network administrators, but also of very particular links amongst these elements – both human and artifact – and of examining the factors that contribute to its overall cohesion and continued ability to function effectively. In essence, these are issues of social order, or rather, socio-technical order, in this system. I propose that socio-technical order is maintained spatially and temporally through assumptions about cultural norms – particularly trust – that find material form in the technology of inter-domain routing.

This begs the question as to what it might be that promotes cohesion and order across this distributed socio-technical system. I believe that we may find the answer by looking to a common culture of network administration. But what, then, is culture in this context? I follow Knorr-Cetina's definition of culture as “the aggregate patterns and dynamics that are on display in expert practice and that vary in different settings of expertise” (Knorr Cetina 1999, 8). She goes on to suggest that culture may be studied as it manifests in material practices, and symbolic structurings, indicating how the “notion of practice shifts the focus away from mental objects such as the interests or intentions that inform concepts of action, and toward the reordered conditions and dynamics of the chains of action of collective life” (Knorr Cetina 1999, 10). I pay particular attention to practice in this work, examining how the culture of network administration is constituted in shared practice around the technology of inter-domain routing.

Knorr Cetina is concerned with cultures of knowledge production, which she terms *epistemic cultures*, drawing back from studying the products of a knowledge process, to studying the mechanics of the process itself. I am focusing on similar processes: those of coordination and communication amongst network administrators to maintain the operational integrity of the Internet's inter-domain routing system. While these are not processes of knowledge making in the sense of scientific research, as are those which Knorr-Cetina studied, there are parallels in the generation of operational knowledge critical to maintaining the stability of inter-domain routing, produced through shared practice. Since my concern is with the mode of functioning of a specific technology – inter-domain routing – rather than the more general social phenomenon of the epistemic culture, I differ from Knorr-Cetina in my explicit focus on the qualitative aspects of the technology itself as a product of a particular epistemic culture: computer network researchers in the United States of America in the late 1980s.

Furthermore, I suggest that the form and widespread deployment of Internet technologies have had material effects on this epistemic culture, especially in the phenomenon of “lock in”. Once widely adopted, particularly in the form of infrastructure, technologies become relatively invariant for two reasons. First, it can be expensive to replace technological deployments in physical infrastructure. Second, and more importantly for this discussion, is the issue of technology standards in markets that exhibit network effects: since all implementations of a technology must conform to the same technology standard in order to interoperate, any major revision to a standard requires coordinated action across all implementations, which is often not a practical choice. As Arthur (1989) famously illustrated, technologies can become “locked in” through widespread adoption. The capacity for technological change can be hobbled by the very success of a technology.

As technologies become “locked in”, features of the technology that are informed by norms prevalent at the time of its design are similarly locked in. These features, in turn, suggest certain forms of practice, which may promote the norms of the culture that originated the technology. The relative invariance of “locked in” technologies is of special significance, as it enables norms from one place and time – the cultural context of technology production – to be made concrete in the form of the technology, and spread spatially and temporally through the technology itself. This is, of course, a slippery causal chain, that is historically contingent within its originating culture, and also contingent on the process of adoption of the technology as it enters new spaces. Nonetheless, the cultural context of the production of a technology remains significant, if only to explore shifting notions of practice, and associated norms, as a technology spreads.

Technology, therefore, is an important element in the study of culture, acting as a possible vector for the propagation of culture through the embedding of norms, alongside more conventional cultural markers such as oral and textual histories.

3.2 Methodological Issues

The notion of the epistemic culture assumes that practices may be studied *in situ*, in the bounded spaces of laboratory settings. This is, however, a problem for the study of network fields, as the setting to be studied does not consist of bounded, or even unbounded, physical spaces. Knorr Cetina (2007) introduces the concept of macro-epistemics to deal with this problem of a lack of place, suggesting that decentred settings of knowledge – ranging from open source software to the global financial system – should also be studied as epistemic cultures. Just as epistemic cultures can be studied through an examination of the material practices at play, so too can macro-epistemics such as the Internet's inter-domain routing system.

Suchman et al (1999), like Knorr Cetina, advise a practice-oriented approach to ethnographic investigations, but with a focus on technology, rather than culture. They are concerned with two reconstructions of technology as practice: first, in the anthropological sense of making sense of the social and cultural worlds of technological systems, and secondly, aimed at creating alternative approaches to technology production. The former is most relevant to the research presented here, as it indicates the useful viewpoint of practice as a constituent element of technology, as well as of culture:

> The question is how to conduct ethnographies of work and technology, including both practices of design and artifacts-in-use, that are aimed at recovering the projects, identities, and interests that inform these practices. (Suchman et al. 1999, 392)

These practices do not, of course, take place in a vacuum; in a very real sense, they produce the spaces in which they occur. A network field may lack bounded physical locations, but this is not to say that it lacks spatiality. In this context, space is constructed as a socially and technically produced quantity, manifesting in the social relations amongst network administrators, and the technological interconnections amongst networks. The useful conception of *network space* has been offered to engage with the investigation of phenomena in this kind of non-Euclidean spatiality, suggesting that “spatiality is an aspect of network stability” (Law and Mol 2001). In essence, a stable heterogeneous network, such as that of inter-domain routing, produces a network space that acts as the field of investigation. In contrast to this perspective, however, I offer a conception that is not entirely non-Euclidean, but rather a combination of Euclidean and network space. The autonomous systems of the Internet, after all, are interlinked systems of physical infrastructure, existing within and across national boundaries. To study this space is to envision it as grounded in physical and political geography, from which more abstract conceptions of socially produced space emerge. This spatial conception is not a binary opposition between Euclidean and non-Euclidean, but rather based on social actors' own understandings of their lived realities.

Even with this conception of space in a network field, my fieldsite could too easily come to be all-encompassing, demanding a full-blown ethnography of network administration. As I had to remind myself several times in the course of this work, my research is a functional investigation of a particular socio-technical system, an approach which allowed me to avoid lines of investigation that were not immediately pertinent to my research questions. This is not to say that I avoided such seemingly tangential investigations, but rather that I chose them with care, and followed them, for the most part, when they fell within these bounds. As such, my research questions acted as additional bounds to my fieldsite, in the selection of investigative paths that I chose to follow.

The choice of geography for this research offers a final set of bounds for my fieldsite. Investigations were, by and large, limited to the United States of America. This was an obvious choice to examine the development of the Internet itself, since most of this development historically occurred in the USA, and was funded by the US government. The production and reproduction of the culture of network administration is not similarly geographically bounded. However, I chose to restrict my investigations to the USA, both for reasons of practicality, and also to gain a richer understanding of the unfolding of this culture in the place of its origination.

Questions remain as to the appropriate means to investigate a network field, and to the validity of the knowledge that might be gained from any such investigation. Ethnographers have struggled with similar issues when faced with local social systems that have global connections, and with systems that are inherently global in nature. Multi-sited ethnography is an approach that engages with these kinds of sites, and which has gained currency in several fields, including science and technology studies and cultural studies of the world system:

> Multi-sited research is designed around chains, paths, threads, conjunctions, or juxtapositions of locations in which the ethnographer establishes some form of literal, physical presence, with an explicit, posited logic of association or connection among sites that in fact defines the argument of the ethnography. (Marcus 1995)

Marcus offers several strategies for multi-sited research, in following people, things, metaphors, stories, biographies and conflicts. Burrell (forthcoming) proposes a conception of the “fieldsite as a network”, to foreground the contours of the phenomenon under study from its larger context. Her strategies were particularly useful to my research process, including advice to seek entry points rather than sites, and to consider multiple types of networks (including telecommunications networks). These approaches still present a difficulty for a network field: they require the ethnographer to be situated at one or more bounded physical field sites and follow links from these locations. In a network field, there is no “place” at which the ethnographer can be situated. Other researchers have dealt with “virtual ethnographies”, investigations into the social worlds that are online, physically placeless, settings. Although these do engage with a setting that is virtual rather than physical, these settings too are bounded and locatable: a researcher can still “enter” and “exit” the virtual reality of Second Life (Hine 2000).

The closest approximations to bounded, locatable settings in my research are represented by conferences and meetings of those involved with network administration, and by email lists frequented by these communities. However, the ongoing practice of managing the inter-domain routing system is not so easily observed, as there is no single locus of coordination. At the most abstract level, one might say that the Internet itself is the object under study; but as my description of the inter-domain routing system shows, the Internet is not a unitary object. Rather, it is a highly distributed graph of socio-technical relationships across autonomous systems, with the administrators of each autonomous system most concerned with the task of coordination with their immediate neighbours. There is no “view from nowhere” (Haraway 1988) that allows us to capture the “Internet”, a fact that network administrators are only too aware of: all perspectives are partial, situated in relations – both spatial and temporal – amongst very specific sets of administrators and networks. Indeed, the question of how much can be known about the Internet's topology, and with what veracity, is an ongoing topic of research in the computer networking community (Oliveira et al. 2008). The only real universal in this system is the technology of inter-domain routing itself, which perforce must be standardised across all interconnected autonomous systems for the Internet to exist at all.

A final useful research strategy follows from Latour's advice to study social groups as ongoing dynamic processes, rather than static configurations. He suggests attention to controversies, both big and small, claiming that we are better able “to find order *after* having let actors deploy the full range of controversies in which they are immersed” (Latour 2005, 23). Controversies disturb network space, and in doing so, surface topologies that may otherwise have remained invisible. In fact, it was the controversy over the hijack of YouTube's address space by the Pakistan Internet Exchange that instigated this research. I found great gains in following controversies, and in deploying them as objects of discussion in my interviews.

I supplement the research strategies presented here with additional modes of observation particular to the phenomenon that I explore: careful attention to the qualities of links amongst the network of people and artifacts constituting the process of inter-domain routing, and to the particulars of the technology itself.

Each of these strategies offers a window onto my fieldsite; taken together, they offer the means to construct a rich picture of the complex world of network administrators involved with maintaining the Internet's inter-domain routing system.

3.3 Materials

The research presented here is based on a range of different materials. Primary textual sources include the North American Network Operators Group (NANOG) email list archives, videos of presentations at NANOG meetings, details of various incidents related to inter-domain routing, papers on inter-domain routing from the networking research community, and standards documents from the Internet Engineering Task Force (IETF). Secondary textual sources include Abbate's (1999) well-told history of the development of the Internet, as well as several other accounts of significant events in the history of the Internet.

In addition, I conducted 18 interviews to gain a first-person understanding of the world of network administration. Interviewees were mostly network administrators, but also included a few individuals who are involved in the design and deployment of inter-domain routing equipment, or who are active in the standards process at the IETF. They ranged in experience from some who were involved with the administration of NSFNET – the Internet's immediate ancestor – in the late 1980s, to others who began their careers as little as five years ago. Interviews were conducted by telephone, in person at my interviewees' offices, and also at the 74th IETF meeting in San Francisco. Interviews lasted from about 30 minutes to over an hour. Contact was made with interviewees through a combination of personal introductions and cold calls. In several cases, a strategy of snowballing was adopted, where interviewees provided introductions to acquaintances as possible interview candidates. In keeping with standard research practice, I have kept the identities of my interviewees anonymous.

My own experience and preparation for this work themselves could be considered tertiary materials. For this investigation, I drew on a decade of experience as a technologist in the software industry, graduate classes on computer networking, and research projects involving the technology of inter-domain routing. While this degree of experience may not be absolutely necessary for such an investigation, I suggest that any exploration of an intensely technological lifeworld requires preparation that engages deeply with the specifics of the technology itself. As Nissenbaum (2001) puts it:

> Sometimes a fine-grained understanding of systems – even down to gritty details of architecture, algorithm, code, and possibly the underlying physical characteristics – plays an essential part in describing and explaining the social, ethical, and political dimensions of new information technologies.

Even with this preparation, I could not possibly have understood this lifeworld without the expertise of my interviewees, each of whom generously gave of their time to make this investigation possible.

3.4 Cyborgs, Deconstructions and Reconstructions

Haraway (1991) presents the compelling image of the cyborg, a view of human and machine as inextricably linked into a single organism, while at the same time calling to attention the politics implicit in shifting human/machine boundaries and reconfigurations. I take Haraway's advice in the sections that follow, first deconstructing the organisation of networks and network administrators to examine spatial form and shifts in this configuration over time, and then exploring the ways in which networks and network administrators function as a seamless whole in the operational context of the Internet, through the mediating social value of trust. In this instance, the power of the cyborg image lies in collective, rather than individual, form.

4 Trust in Technology

In this section, I present an account of the evolution of inter-domain routing, and the manner in which trust came to be such an important element of its constitution. This is at once a story of technological progression and deployment, and of the emergence of distinct social groups and institutions. The analysis presented here draws from the Social Construction of Technology (SCOT) approach (Pinch and Bijker 1987), which provides a useful framework for picking apart these complexities. SCOT calls to our attention the social groups that are relevant to decisions taken around a technology, and the particular technological frames through which actors in these groups engage with one another. In addition, SCOT suggests that technologies stabilise after a period of negotiations amongst relevant social groups, leading to closure around a stable form of the technology. It is possible that each group may not view the technology the same way, even in its stabilised form, a phenomenon SCOT refers to as interpretive flexibility. These elements are all visible in the period that I consider here, covering two decades since the creation of BGP in 1989.

Before the Internet came NSFNET, a non-commercial network operating from 1986 to 1995, within the United States of America, which served as the test bed for the development of Internet technologies, and also provided connectivity for a number of research institutions. The NSFNET, in turn, was preceded by ARPANET in the 1970s, a network funded by the Advanced Research Projects Agency (ARPA) of the United State's Department of Defence. The atmosphere amongst the early ARPANET researchers was collegial, rather than formal, which also had the effect of increasing the sense of involvement and commitment amongst researchers. As Carr, Crocker and Cerf 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, Crocker & Cerf 1970 quoted in Abbate 1999).

These qualities of informality and trust amongst early computer networking researchers contributed much to the development of the ARPANET. They also – very importantly – laid the foundations for the evolution of social institutions of the NSFNET, in which context BGP was developed. As one of my interviewees, who was involved with NSFNET in the late 1980s, commented:

> In the early days, the atmosphere was really, really different. People were very cooperative. We're all thinking that we're doing a great thing here, it's historical. People just wanted to do good.

BGP was developed under the aegis of the Internet Engineering Task Force (IETF), a standards body which has no statutory standing, but which nonetheless is responsible for the development of the Internet protocols. The IETF follows the informal arrangements of the ARPANET era; there is no formal process to becoming a member of the IETF, it is sufficient to contribute ideas on an IETF mail list, or at an IETF meeting.

BGP version 1 was standardised by the IETF in 1989, to supersede its predecessor, the Exterior Gateway Protocol (EGP). The principal reason for this shift was to remedy the load imposed by EGP on the NSFNET. EGP routers sent periodic updates of their routes every 3 minutes, an approach which performed reasonably well when there were only a few hundred routes on the NSFNET. However, as the number of routes increased, this update process began to cause a noticeable reduction in network performance, which was the immediate driver for the creation of BGP.

BGP is sometimes known as the “Three Napkins Protocol”, so named for the fact that it was initially sketched out at an IETF meeting on three paper napkins by Yakov Rekhter, Len Bosack and Kirk Lougheed. It took less than a month to create two interoperable implementations of BGP version 1 after this first specification was written. BGP was intended to be a short to medium term solution for inter-domain routing on the NSFNET; its authors did not anticipate that it would come to be the de facto protocol for inter-domain routing on the Internet (White, McPherson, and Sangli 2004).

One could say that there were two social groups involved in negotiating this shift from EGP to BGP – the IETF, and the network administrators of the time – but in truth, they were practically one and the same. This quote from one of my interviews illustrates the relative ease in the process of standardisation in the NSFNET era, and suggests how things have changed over the years:

> If you dig up the [BGP version 1] Internet draft, you can see the date of the first Internet draft, it was then approximately that we were doing the work. It's not like nowadays, you know, at that time we would just write the document and start to implement it and deploy it. It's not like today, you have to debate. I don't know if you're familiar with the IETF? If you have a new protocol, it's going to be debated for a long time, because there are a lot of commercial interests, companies' interests, behind it.

NSFNET was structured as a hierarchy, with the NSFNET backbone at the top, followed by a second tier of regional networks, and finally institutional networks. The NSFNET backbone reached across the United States, connecting 30 regional networks, which in turn provided connectivity to over 200 academic and research institutions. For instance, UC Berkeley was connected to the Bay Area Regional Research Network (BARRNET), which in turn was connected to the NSFNET backbone. The significance of this structure lies in the fact that the NSFNET backbone acted as a central point of control, quite unlike the mesh of interconnections on the Internet of today.

Since NSFNET only allowed non-commercial traffic, a number of regional networks arose to support commercial demand. Recognising this shift, the NSF issued a solicitation in 1993 for the creation of the next generation of the NSFNET (NSF 1993), to serve both commercial and non-commercial traffic. This resulted in the privatisation of the NSFNET in 1995 to form the Internet. The North American Network Operators' Group (NANOG) was created in this period, initially funded by the NSF, to serve as a means for coordination amongst network operators, and also as a forum for planning the transition from NSFNET to the Internet. At the same time, companies building Internet routing hardware, such as Cisco, became increasingly involved in the IETF, since they would be the ones actually implementing protocols such as BGP in their products. For instance, engineers at Cisco, including Kirk Lougheed and Tony Li, co-edited every version of the BGP standard in this period leading up to BGP version 4 in 1995.

We see, therefore, a shift from a single community to a situation where at least three distinct social groups are involved in standardisation efforts for BGP: vendors of routing hardware, network administrators, exemplified by NANOG, and networking researchers at the IETF. However, it would be naïve to treat these as completely independent groups in this period; the same people were often involved in multiple settings, and most of those involved came out of the tightly knit NSFNET research community. Even so, it is useful to consider these gradual shifts as part of a process that led to more clearly demarcated separations amongst these social groups over time.

A series of BGP specifications were standardised in the NSFNET era, from version 1 in 1989 to version 4 in 1995, which is currently in use on the Internet. The stabilisation of BGP, therefore, occurred over a period of time during which the social groups involved were themselves in a process of formation. The close ties amongst these social groups resulted in closure on a form of BGP that was "trusting": mechanisms to validate route advertisements were never built into the protocol. In my interviews, a common refrain has been that the trust engendered by the close ties amongst the "Internet community" were a significant factor in this lack of security. Since network administrators often knew one another personally, and trusted their peers to make sensible BGP configurations, there was no need anticipated for security in the protocol.

In addition, the hierarchical structure of the NSFNET allowed the NSFNET backbone administrators control over the routes that NSFNET would carry, ensuring that false BGP advertisements from any regional or institutional network would be stopped at the backbone, and not propagated across the NSFNET as a whole. Whenever a new block of IP address space was allocated to any of the networks on NSFNET, an email exchange would occur between the administrator responsible for the regional network which would be routing this address space, and the administrators at the NSFNET backbone, requesting that this new route be entered into the list of routes allowed by the backbone. These routes were maintained in a repository administered at the NSFNET backbone, called the Policy Routing Database (PRDB). As an administrator involved with operating the NSFNET backbone told me:

> What we did was a lot of manual work. We had regional designated persons, we have certain names... So 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.

In essence, a combination of a hierarchical network structure, and a community of trusted individuals, allowed the NSFNET to maintain routing integrity, with no need for security mechanisms to validate BGP advertisements. However, this does not account for why security was not a principal concern in the design of BGP. One might imagine, in retrospect, that a protocol intended for the political purpose of separating network administrations should be secured, that it should not be quite so trusting.

In fact, the trust amongst the members of this tightly knit community played a central role in this decision, a theme which many of my interviewees commented on, for instance:

> The security issue was one that was not addressed initially, because it was not a commercial Internet, everybody trusted everybody, we were all hackers.
>
> Security was not a big issue at that time, people just want to make things work.
>
> 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.

Although numerous extensions have been added to BGP since 1995, there was closure around the basic mechanisms of the protocol with BGP version 4. There was not much interpretive flexibility in this process; as these social groups were quite tightly knit, they had a common interpretation of BGP at this time. BGP was created in the image of the social context of its production – “everybody trusted everybody” – a legacy that remains with us to this day.

As Shapin (1995) has illustrated, trust is an essential element of the research enterprise, and indeed, of epistemic machinery in general. If we accept that the creation of new knowledge is contingent upon the existence of prior bodies of knowledge, then the very act of knowing is also an act of reposing trust in those who created these prior knowledges. With this in mind, it is perhaps not so surprising that BGP is as trusting as it is, given that it was created in the operational context of a research community.

After the privatisation of NSFNET in 1995 to form the Internet, the emerging social groups of equipment vendors, network administrators and network researchers separated much more clearly from one another. It is not an uncommon occurrence today, for instance, to see complaints on the NANOG email list that the IETF does not pay attention to the concerns of network operators. Equally, network operators have their own view of vendors as being disconnected from operational concerns. When I mentioned to one of my interviewees, a network administrator, that I was also planning to talk to engineers at Cisco, he warned me to keep in mind that while vendors were adept at protocol design, they had very little idea of day-to-day operational issues. Another interviewee remarked on the disconnect between network administrators and “seriously top notch” networking researchers at a meeting he went to, telling me how the researchers had little idea of the practice of network administration. Influential voices from the early Internet community have gone so far as to publicly call the IETF the “IVTF”: the Internet *Vendor* Task Force (Bush 2005).

After 1995, the Internet spread across the world at a fairly explosive rate, in ways that those handling the NSFNET to Internet transition did not entirely anticipate. The NSF's 1993 solicitation called for three distinct components for the Internet: Network Access Points (NAPs), a Routing Arbiter (RA) and very high speed Backbone Network Services (vBNS). The NAPs were intended to be points at which networks – both commercial and non-commercial – could interconnect with one another, and vBNS was to continue in the tradition of the NSFNET, providing high speed connectivity to academic and research institutions. The most important of these three, for the purposes of this investigation, was the Routing Arbiter:

> 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)

The Routing Arbiter, in essence, was to take on the role of the NSFNET backbone as a trusted repository of authoritative routing information. It was envisioned that network operators would maintain their routing information in this repository, and use information provided by other operators to construct their routing policies. The critical difference from the NSFNET backbone is that the RA was to be a third party service, rather than an integral component of the Internet's structure. As the solicitation itself recognised, the RA "*may* be used by attached networks"; it did not *have* to be used. Indeed, there was no way to mandate usage, since the RA was not itself responsible for carrying data traffic, unlike the NSFNET backbone.

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)[6], which remains in use today; a shared, publicly available, store of routing information, with this information maintained voluntarily by network administrations. The European Internet registry, RIPE, created a similar database, and their database format, RIPE-181, was adopted as a standard for 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 organisation that took over the NSFNET backbone, Advanced Network and Services (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

[6] See http://www.ra.net

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. 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.”

The initial incarnation of RADb had several problems, some of which have since been rectified. Dates associated with updates to RADb records were originally maintained manually, as a result of which the most current update to a route was not necessarily the one with the most recent date. Anybody could enter routing information, leading to situations where administrators of a network would sometimes maintain data for their customers' networks in RADb, so that they would be able to use RADb to generate their own routing configurations. When these commercial relationships changed, the provenance of this data was, more often than not, lost. Worst of all, routing information was put into RADb all too readily to ensure routing by ANS and other large network operators, but outdated information was never taken out, making it difficult to distinguish the wheat from the chaff.

Some of the major new backbone networks, including Sprint, AT&T and UUNET, came to view information about their interconnections with other networks as commercial secrets, and ceased to publish updated routing information to the IRRs. As one of my interviewees who was involved with RADb noted:

> ... some of the ISPs, because of competitiveness, they don't really want other people to know who

> their clients are. If you update RADb, it's public, you say, I'm going to announce this prefix, and you know, Sprint has whatever XYZ as their customer, and for a period of time, I don't know why, maybe it's competition, a lot of ISPs like to keep it secret, who their clients are, so that doesn't give them incentive to update RADb, and they just try to grow their own database.

In addition, as the Internet service provider industry has moved through a series of mergers and acquisitions, and as new administrators come on board to manage networks, the information in the IRRs has come to be inconsistent in terms of quality. In spite of these problems, several large provider networks still require their customer networks to maintain routing information in the IRRs. As these developments illustrate, a fourth social group emerged after 1995, that of the firms which own and operate networks, the network operators.

Matters become still more complex once the process of assigning IP address prefixes to networks is considered. Some of the original records of prefix allocations have been lost, as participants in one of the sessions at the 74th IETF noted, making it difficult to determine who actually has the right to originate these prefixes. In addition, networks may delegate portions of their prefixes to other networks, notifying IANA, or the appropriate RIR[7], of this delegation. One of my interviewees, an administrator of a large research network, related how he had delegated a portion of a prefix to a departmental office. That office has long since ceased to exist, but it is difficult (although not impossible, in this particular case) to provide evidence to retrieve the delegated address space.

With all these uncertainties, administrators are understandably careful about customers'

[7] IANA – Internet Assigned Numbers Authority, RIR – Regional Internet Registry. See Section 2 for details.

requests to route a new block of address space. This is a relatively straightforward process if the customer obtained the address space directly from the network in question. However, customers may also obtain address space from another network, or directly from a RIR. An administrator at a Tier 1 network described the processes in his organisation for checking such requests before modifying BGP configurations to route the requested address space:

> ... we're usually pretty diligent about checking the SWIP[8] 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.

Even with the checks and balances built into this process, it can be difficult to verify the authority of the party requesting the change. Subtle challenges arise when trying to resolve whether or not an individual at a company – which does have rights to a certain address block – actually has the authority to request changes to the routing of that address space. More worrisome is the issue of fraud: individuals set up companies with names very similar to defunct companies to try and steal address space registered to these companies. One such story came up in my interviews, illustrating how these challenges have changed the administrative practices around managing IP address blocks:

> They found somebody who's not actively using some space, incorporated in a different state under a very similar, perhaps even identical, name, got the corporate paperwork, generated a letterhead,

[8] Shared WHOIS Project, a mechanism for maintaining and checking records of IP address block assignments.

> sent in a letter to ARIN[9] saying, we want to update our registry of this legacy space from way back when; we moved, we're not in Colorado... and all of a sudden that address block becomes a big home for spammers. So ARIN now is a *very* careful about doing those transfers. They not only want a statement showing somebody of the same name exists, they want a legacy of how it got from where it was to where it is. And that's proving very difficult, because almost nobody is around any more who remembers it.

As this history of shifting practices illustrates, the interpretation of BGP changed from the time of its creation to the Internet of today. The social groups involved in standardising BGP were tightly knit and had a common interpretation of it, creating a "trusting" protocol. Network operators, on the other hand, interpret BGP as an expression of commercial relationships of network interconnection, and are therefore less willing to entrust their BGP configurations to the public domain. At the same time, however, there are many network administrators, particularly those from the early days of the Internet, who view the actions of these network operators as a violation of the expectations of disclosure and trust implicit in the Internet community. Not only do these social groups have different interpretations of BGP; they also perform these interpretations in distinctly different technological frames.

The upshot of these developments is that neither BGP, nor the IRRs, provide a reliable mechanism for validating advertisements. This has increasingly become a cause for concern over the last decade or so. As the Internet spread on a global scale, it has become difficult to maintain the close ties of the NSFNET period. Errant, and even malicious, BGP configurations have resulted in instabilities on the network. As a result, it seems apparent that the network

[9] The Regional Internet Registry for North America.

administrators', vendors' and the IETF's technological frames around BGP have shifted from viewing it as part of a socio-technical system of which trust in an important element, to one which anticipates deception.

This loss of trust has resulted in a renegotiation of the practice of inter-domain routing on the Internet, as was made concrete in BGP. For instance, filtering of customers' advertisements is now much more common that it was in the late 1990s. This is perhaps amongst the principal reasons for the continued stability of the Internet as a whole. However, it remains impractical to apply filters between Tier 1 networks, and the larger Tier 2 networks[10], an issue to which I return in the next section.

[10] See Section 2 for details.

5 Trust in People

In this section, I focus on the social world of network administration, considering the technology of inter-domain routing as a necessary substrate to these social relations. I analyse it as historically contingent on the trusting culture of the NSFNET period, since many of those involved with the NSFNET went on to work in commercial network providers. At the same time, I explore how the trust implicit in BGP has continued to shape practices of inter-domain routing. I contend that these two factors together have contributed towards the continuance of a common culture of network administration that demands a degree of trust amongst its constituents.

I was surprised by how easily words like “community”, “family” and “trust” came up in my interviews. These are, to network administrators, concepts that are deployed as part of everyday practice, and are used to claim membership in a whole that is greater than the sum of the individual networks which they administer. It is easy to understand why this was so in the NSFNET period, considering how that was a relatively small research community. As the NSFNET transitioned to the Internet, however, these conditions do not hold true any more: the community that does exist is neither small, nor research-oriented. In part, the sense of community is a progression from the NSFNET period, since many of those involved in the NSFNET moved on to positions in various commercial networks, such as ANS, carrying the social ties of the NSFNET community with them. As an administrator who was involved with the NSFNET 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.

The NSFNET personnel brought with them the practices of the NSFNET, and inculcated these into the next generation of network administrators, both through direct experience, and through simple operational demands. For instance, the requirement from ANS, discussed in the last section, that all networks routing through its infrastructure had to register their routes with RADb. At the same time, a great deal of circulation of personnel amongst network operators was in progress, further generating social ties that spanned individual network operators. One of my interviewees recalled that period:

> ... there was a time, not long ago, when if you found anybody who had been with a given network organisation 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!

Just as it is useful to understand a community in terms of inclusion, it is also useful to ask who is excluded from a community, to get a better sense of where the boundaries are drawn. Another one of my interviewees drew a sharp distinction between the managers running the business side of networks, and the network administration community:

> That's how it worked, a lot of it was just of a cooperative nature. While I said, the Internet is not all one big happy family today, and it isn't, but for the most part it is. At the technical level it very much is. We get along amazingly well at the technical level. The corporate sides, like the depeering between Sprint and Cogent, or the Cogent-Level 3 depeering[11] 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 realise what havoc they were going to wreak for their *own business*. They weren't just causing other people problems, they were causing their own people problems.

As these quotes suggest, my interviews indicate a cohesive social world that extends across the networks – commercial and non-commercial – that make up the Internet. I began my research in search of this community, expecting it to be relatively undifferentiated, and mediated through the NANOG email list and meetings. I could not have been more wrong; this is a complex and fascinating social world, structured by reputation and the relative positioning of networks in the Internet.

For a network's administrator to have a sense of being part of this community, it is not enough for a network just to be connected to the Internet. The network must be of sufficient size and complexity for inter-domain routing to be an issue that is important to its administrators. Administrators at small ISPs that I talked to had no sense of being part of the larger Internet community, since they connected to the Internet through a single provider network; their

[11] Incidents where two networks disconnect from one another. These were of particular significance since the size of the networks involved resulted in a partition of the Internet: some of these networks' customers could not reach one another. See http://www.renesys.com/blog/2008/10/wrestling-with-the-zombie-spri.shtml for more on the Sprint-Cogent dispute, and http://www.forbes.com/forbes/2008/1013/064.html for a history of Cogent.

relatively small size does not justify having redundant connections through multiple providers. For them, BGP configurations are something that the networks which they connect to, provide to them. They may tweak these configurations, but they have no choice about their path to the Internet. As Luhmann (2000) puts it, since they lack choice, they also lack risk in the everyday practice of inter-domain routing; they have confidence, rather than trust, in their provider network. In this context, at least, community and trust are inextricably intertwined concepts: to lack trust is also to lack a sense of community.

This is a dynamic that may shift over time, as a need is felt for a network to have connections to multiple peers. A network administrator at a university campus, who has been involved with managing networks since the late 1980s, had this to say of transitions from the NSFNET period:

> Because we connected to BARRNET[12] 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 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.

As Internet connectivity became of greater importance to the campus, it acquired multiple peers, interconnecting directly with other campuses across California, and through commercial and non-commercial networks to the larger Internet. This multiplicity of connections brought with it a choice of paths to different destinations on the Internet, and risk associated with making these choices, creating social relations of trust that follow the interconnections of

[12] Bay Area Regional Research Network, one of the NSFNET's regional networks.

networks. Administrators in each pair of peer networks trust that their counterparts will maintain their interconnection responsibly. This is not, of course, a blind trust; connections are actively monitored, and choices are revisited based on the perceived quality of connections. These are serious issues for any network, and are magnified as the number of peers, and therefore the range of choices for routing, increases: Tier 1 and Tier 2 networks may have peers numbering in the hundreds, or even the thousands.

The trust embedded in BGP, however, is not just between pairs of peer networks; it is transitive across all the interconnected networks of the Internet[13]. To trust a peer network is to trust all of its peers, and so on, generating a web of trust that spans the Internet. The social relations of trust between network administrators in pairs of peer networks are therefore extended across the Internet through the particular form of the technology of inter-domain routing. It is this feature of inter-domain routing – transitive trust – which, I suggest, provides the technological underpinning for a sense of a community operating at the scale of the Internet itself; especially including all those involved with the administration of Tier 1 and Tier 2 networks.

A famous incident of a failure of these relations, which illustrates the extended dependencies, is the case of AS7007, a story which has entered the collective lore of the network administration community[14]. On 25th April, 1997, AS7007 began announcing advertisements for an immense number of prefixes, most of which it was not capable of routing. One of its

[13] See Section 2 for a detailed discussion of transitive trust in BGP.

[14] See http://flix.flirble.org/ for more details about this incident. The explanation and apology from the administrator responsible for AS7007 to the members of the NANOG email list is available at http://www.merit.edu/mail.archives/nanog/1997-04/msg00444.html.

customers had erroneously announced an unusually large number of prefixes to AS7007. Instead of filtering these advertisements, AS7007 expanded this set of prefixes (through misconfiguration and/or software bugs), and relayed them to the rest of the Internet, modifying the advertisements to claim that *it* was the autonomous system that originated these advertisements, rather than just a relay, thereby removing the means by which other autonomous systems could make judgements about the veracity of these advertisements. BGP routers around the world began selecting AS7007 as a preferred route to almost every prefix on the Internet, instead of the networks which were actually capable of routing these prefixes. The number of prefixes announced was so large, in fact, that many routers ran out of memory and crashed. Although this was corrected within a period of a few hours of the original fault, the result was a global failure of the inter-domain routing infrastructure.

This incident, along with others like it, was instrumental in shifting the practice of inter-domain routing towards increased filtering of customers' BGP advertisements. However, it is less practical for Tier 1 and Tier 2 networks to apply filters to their connections with one another, both due to the sheer volume of advertisements on these connections, and also because it is difficult to evaluate the veracity of advertisements when an autonomous system is only acting as a relay for advertisements originating elsewhere. An administrator at a Tier 1 network described these difficulties to me:

> 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 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 authorised to advertise this prefix, but when it's through a peer, that's a little bit harder.

In the core of the Internet, therefore, there can be no substitute for trust amongst the administrators of these networks. This is reflected in the attendance at NANOG meetings: one of my interviewees commented that attendees generally tend to be personnel from Tier 1, Tier 2 and academic networks. When I asked an administrator at a Tier 1 network how important he thought the relationships formed at NANOG were to maintaining the global inter-domain routing system, he responded:

> Oh, very, very important, very important. 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.

Gaining admission to this community can be a difficult process. A network administrator, who is also an active researcher, told me of the challenges he faced when trying to set up a network monitoring service in the late 1990s:

> ... 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 ... 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.

The trust in these relationships is not a binary quantity; the question is not simply whether to trust, or not to trust, but how to trust, in what degree, and with what resources. Trust, therefore, is contingent on reputation, which is captured in part through the concept of “clue” that the network administration community uses as an expression of expertise. To be “clueful” is to have a high degree of expertise. In talking about practices of filtering, for instance, a Tier 1 network administrator told me that, “on the peer side today, we more or less trust our peers to have a high enough clue factor, and for them to filter their customers, and for them to do it well.” Clue, or lack of clue, were recurring themes in my interviews, and on the NANOG email list, where it is not an uncommon occurrence to find requests for contact with “clueful engineers” from one network or another to help resolve a problem. The “Cisco Clue” wiki[15] declares: “Have Cisco routers? Have clue or need a clue? If so, this may be the site for you.”

To be clueful is necessary, but not sufficient, to have a good reputation in the community. An administrator must also participate in the life of the community, demonstrate cluefulness, and maintain a network which others know to offer a high quality of service. One of my interviewees described what it might take to enter the community as a “trusted individual”:

> ... the only way that you gain admission into a community is to be a trusted individual. 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.

[15] See http://cisco.cluepon.net

In addition, reputation is based on the relative position of a network in the Internet. When I asked an administrator at a Tier 1 network how important he thought it was to have stature in the NANOG community, he 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.

It is a matter of prestige in the network administration community to be an administrator at a Tier 1 network. One of my interviewees operates a network that is, for all intents and purposes, a Tier 1; however, he is concerned for the day that some of his network's peering agreements might be renegotiated, “We worry about it, we want to keep our status and not have any [paid] transit.”

As is made clear through these conversations, reputation and associated relationships of trust in the network administration community are structured in a form following the structure of the Internet itself, even as these social relations spread transitively across the Internet.

6 Conclusion

There is a debate amongst scholars around the place of trust in understanding society. Some believe that trust is but one element of social reality, while others contend that it is the very element that allows any kind of social – or socio-technical – order to exist (Gellner 2000, 142). My research suggests that, at least in the context of the Internet's inter-domain routing system, trust is coextensive with the socio-technical order. Gellner deals with the problem of social order and trust, contrasting Hobbes and Ibn Khaldun:

> The Hobbesian problem arises from the assumption that anarchy, absence of enforcement, leads to distrust and social disintegration. We are all familiar with the deductive model which sustains and re-enforces that argument, but there is a certain amount of interesting empirical evidence which points the other way. The paradox is: 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. This is a basic fact about the human condition, or at any rate about a certain range of real human conditions. It is the basic premise of Ibn Khaldun's sociology[...]. (Gellner 2000, 143)

Even though Ibn Khaldun's sociology referred to 14th century Arab society, it is fascinating to think that the inter-domain routing system may well fall within the range of real human conditions that Gellner suggests would fit this sociology. As I have illustrated, inter-domain routing is a highly decentralized system, with no formal governing body of any real power. Autonomous systems and their administrators are beholden to their peers, rather than to any central authority. A recent email on the NANOG list says it better than I possibly could:

> The Internet's greatest strength and greatest weakness is the lack of a central authority who can "just do it". I for one am happy it is that way. It's part of what makes us an *autonomous* system, sovereign of our own little kingdom.[16]

I extend traditional sociological analyses, to examine the strong mediating presence of technology with embedded social norms. As I have argued, the technology of inter-domain routing (BGP) itself embeds transitive trust, as a consequence of the social context of its development. This transitive trust in the technology, in combination with the social norms of trust from the NSFNET period, have contributed to the sense of community amongst network administrators in the Internet of today. I have also shown how these social relations of trust are structured in this community in a form following the structure of the Internet itself.

There are several possible future directions to move this research forward. One of them is the investigation of instabilities that have recently been occurring at the peripheries of the Internet. In the last year alone, there were reports of faults in the inter-domain routing system in Brazil, Russia and Turkey. Other instabilities have more recently appeared at what might be considered another technological periphery: that of the software implementing inter-domain routing. For instance, a recent incident originating at a Czech network took down a number of routers which were running older versions of software. Another possible direction for research is a broader examination of the cultures of the Internet. For instance, how might this understanding of the the inter-domain routing system influence an analysis of the arguments for and against network neutrality? Could it be said, perhaps, that the network neutrality movement is, in part, an outcome of the trusting culture engendered by these technologies?

[16] From http://www.merit.edu/mail.archives/nanog/msg17353.html

Trust, in this instance, is *the* essential element of social reality. However, this reality is not just social, but also technological; the spatial and temporal propagation of trust cannot be understood without taking into account the particulars of these technological elements. Any study of socio-technical systems must carefully examine the context of production of technologies, and the embedding of social and cultural norms from this context, manifesting in specific material technological practices. This is of particular importance for infrastructures, since it demonstrates how social and cultural norms from one context may come to exert themselves in another context, over great distances and periods of time.

The significance of infrastructure is not simply in the matter of “lock-in” of technologies and particular technological deployments. More important is the issue of how infrastructures become an unconscious part of our everyday lives:

> This powerful infrastructural logic which allows the world to show up as confident and in charge is rarely written about in and for itself [...] and yet this `emptiness' lies at the root of our being, producing senses of the rightness and wrongness of the world so fundamental that we find it difficult to articulate them or to consider that these senses could have been otherwise. (Thrift 2004, 176)

Communications technologies, such as the Internet, are the most widespread infrastructural technologies in the world today, heralded as the driving force behind the new form of the information society. It is critical that we engage with the logic of this global infrastructure, to understand how the ways in which we have reshaped the world have come to shape our societies, and our selves.

7 References

Abbate, Janet. 1999. *Inventing the Internet*. The MIT Press.

Arthur, W Brian. 1989. Competing Technologies, Increasing Returns, and Lock-In by Historical Events. *Economic Journal* 99, no. 394. Economic Journal: 116-31.

Brown, Martin. 2008. Pakistan hijacks YouTube. *Renesys Blog*. http://www.renesys.com/blog/2008/02/pakistan-hijacks-youtube-1.shtml.

Burrell, Jenna. The Fieldsite as a Network: a strategy for locating ethnographic research. *Field Methods* (forthcoming).

Bush, Randy. 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, no. 5: 67-68. doi:10.1145/1096536.1096544.

Edwards, Paul N. 2003. Infrastructure and Modernity: Force, Time, and Social Organization in the History of Sociotechnical Systems. In *Modernity and Technology*, ed. Thomas J. Misa, Philip Brey, and Andrew Feenberg, 185-226.

Gellner, Ernest. 2000. Trust, Cohesion, and the Social Order. In *Trust: Making and Breaking Cooperative Relations*, ed. Diego Gambetta, 142-157. Basil Blackwell.

Haraway, Donna. 1988. Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. *Feminist Studies* 14, no. 3: 599, 575.

—. 1991. A Cyborg Manifesto: Science, Technology, and Socialist-Feminism in the Late Twentieth Century. In *Simians, Cyborgs and Women: The Reinvention of Nature*, 149-181. New York: Routledge. http://www.stanford.edu/dept/HPS/Haraway/CyborgManifesto.html.

Hine, Christine M. 2000. *Virtual Ethnography*. Sage.

Huffaker, Brad, and kc claffy. 2008. Visualizing IPv4 Internet Topology at a Macroscopic Scale. http://www.caida.org/research/topology/as_core_network/.

Huston, Geoff. 2009. CIDR Report. April 18. http://www.cidr-report.org/as2.0/.

Knorr Cetina, Karin. 1999. *Epistemic Cultures: How the Sciences Make Knowledge*. Harvard

University Press.

—. 2007. Culture in Global Knowledge Societies: Knowledge Cultures and Epistemic Cultures. In *The Blackwell Companion to the Sociology of Culture*, ed. Mark D. Jacobs and Nancy Weiss Hanrahan, 65-79. Blackwell.

Latour, Bruno. 2005. *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford University Press, USA.

Law, John. 1992. Notes on the theory of the actor-network: Ordering, strategy, and heterogeneity. *Systemic Practice and Action Research* 5, no. 4: 379-393. doi:10.1007/BF01059830.

Law, John, and Annemarie Mol. 2001. Situating technoscience: an inquiry into spatialities. *Environment and Planning D: Society and Space* 19, no. 5: 609–621. doi:10.1068/d243t.

Lewis, J. David, and Andrew Weigert. 1985. Trust as a Social Reality. *Social Forces* 63, no. 4 (June): 967-985. doi:10.2307/2578601.

Luhmann, Niklas. 2000. Familiarity, Confidence, Trust: Problems and Alternatives. In *Trust: Making and Breaking Cooperative Relations*, ed. Diego Gambetta, 94-107. Basil Blackwell.

Marcus, George E. 1995. Ethnography in/of the World System: The Emergence of Multi-Sited Ethnography. *Annual Review of Anthropology* 24 (October): 95-117.

Nissenbaum, Helen. 2001. How Computer Systems Embody Values. *Computer* 34, no. 3: 120-119.

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. May 6. http://w2.eff.org/Infrastructure/NREN_NSFNET_NPN/nsf_nren.rfp.

Oliveira, Ricardo V., Dan Pei, Walter Willinger, Beichuan Zhang, and Lixia Zhang. 2008. In search of the elusive ground truth: the Internet's AS-level connectivity structure. *SIGMETRICS Perform. Eval. Rev.* 36, no. 1: 217-228. doi:10.1145/1384529.1375482.

Pinch, Trevor, and Wiebe Bijker. 1987. The Social Construction of Facts and Artifacts: Or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other. In *The Social Construction of Technological Systems*, 17-50. MIT Press.

Piscitello, David M., and A. Lyman Chapin. 1993. *Open Systems Networking: TCP/IP and OSI*. Addison-Wesley.

Rekhter, Yakov, and Tony Li. 1995. RFC 1771 - A Border Gateway Protocol 4 (BGP-4). http://www.faqs.org/rfcs/rfc1771.html.

Shapin, Steven. 1995. *A Social History of Truth: Civility and Science in Seventeenth-Century England*. University Of Chicago Press.

Star, Susan Leigh. 1999. The Ethnography of Infrastructure. *American Behavioral Scientist* 43, no. 3: 377-391. doi:10.1177/00027649921955326.

Suchman, Lucy. 2009. Agencies in Technology Design: Feminist Reconfigurations. In *Online Proceedings of the 5th European Symposium on Gender and ICT*. University of Bremen. http://www.informatik.uni-bremen.de/soteg/gict2009/proceedings/GICT2009_Suchman.pdf.

Suchman, Lucy, Jeane Blomberg, Julian E. Orr, and Randall Trigg. 1999. Reconstructing Technologies as Social Practice. *American Behavioral Scientist* 43, no. 3: 392-408. doi:10.1177/00027649921955335.

Thrift, Nigel. 2004. Remembering the technological unconscious by foregrounding knowledges of position. *Environment and Planning D: Society and Space* 22, no. 1: 175–190. doi:10.1068/d321t.

White, Russ, Danny McPherson, and Srihari Sangli. 2004. *Practical BGP*. Addison-Wesley Professional.