# The New Cartographers: Trust and Social Order within the Internet Infrastructure

Ashwin Jacob Mathew
*ashwin@ischool.berkeley.edu*

Coye Cheshire
*coye@ischool.berkeley.edu*

School of Information
University of California, Berkeley

---

*"Order is never observed; it is disorder that attracts attention because it is awkward and intrusive." – Eliphas Levi*

---

## Introduction

Internet governance is often studied in terms of the overlapping interests of nation states and those of formal organizations, such as the Internet Corporation for Assigned Names and Numbers (ICANN) and the Regional Internet Registries, largely in the context of allocation and administration of resources such as Domain Name Server (DNS) designations and Internet Protocol (IP) address space. In addition to recording and allocating resources, governance implies the maintenance of order among those who use these resources. Social order may be maintained through a variety of ways, including the use of legitimate force within a nation state (Weber 1958), market mechanisms or hierarchical command-and-control structures within organizations below the level of the state (see Powell 1990 for a survey). However, the maintenance of order in the infrastructure of the Internet presents a special challenge because the Internet spans all nation states, economies, and organizations. In this paper, we describe a history of Internet governance rooted in social norms of trust. We argue that this informal system has evolved amongst network administrators to maintain stability, and therefore one type of order, in routing flows of data amongst the many individual networks which collectively form the Internet.

To understand the fundamental problem of order in the Internet infrastructure, we must begin with how information gets from one place to another. Routing flows of data between networks is termed inter-domain routing. If we think of the postal service in the United States as an analogy to the Internet's inter-domain routing infrastructure, post boxes would be starting and ending points for information that travels though many different nodes (e.g., postal stations) before reaching an intended recipient. The problem of routing information, whether it is a postal letter or electronic data, involves two questions. First, how does one know what is available on the network? Second, how does one know what path(s) to take to get the information to the recipient when there are no direct paths? The reason that a letter dropped into a mailbox in North Carolina can arrive at the intended recipient's address in San Francisco is largely accomplished because both routing problems have been solved.

The postal service analogy is helpful for thinking about the structure of routing in a large network. However, the analogy fails once we consider the fact that addressing and message delivery are centrally

managed and maintained by the postal system. On the Internet, there is no centralized repository of addressing and routing information. A helpful analogy in this case would be learning a new environment by exchanging hand-made maps with others. If you arrive in a new location and someone gives you a copy of a map with routes through the region, how do you know that you can rely on the information? By extension, can you trust the person who gave you the map? What interest does the person have in giving you a reliable map in the first place? The problem can be infinitely extended if the map is passed from one person to the next over time, extending the range of the mapped territory, and taking into account changes in routes over time. Relying on the routes, geography, topography, etc requires one to depend on every entity that has ever added to the map.

The internal logic of any system is perhaps clearest when it fails, as the Internet's inter-domain routing infrastructure did for YouTube (the popular online video sharing site) on February 24th 2008. For about two hours on that day, the Internet believed YouTube's servers to be located in Pakistan. This kind of confusion seems absurd, but is in fact not an uncommon occurrence (Ju et al. 2010). In this instance, the Government of Pakistan issued an order to censor a particular video on YouTube [link to government order on Global Voices], and the Pakistan Internet Exchange (PIE) chose to implement this by issuing a claim to all networks connecting to the larger Internet through them that they were aware of a more specific route to YouTube, which PIE directed to a dead end, effectively blocking YouTube in Pakistan. Unfortunately, this claim also spread outside Pakistan, through PCCW, an Internet service provider in Hong Kong, which provides international connectivity for PIE. In response, PCCW disconnected PIE until the rest of the Internet could recover from this routing fault. The reason for this failure is that the technology which supports the Internet's inter-domain routing infrastructure – the Border Gateway Protocol (BGP) – provides no mechanisms to evaluate the veracity of claims that networks make to one another about routes which they can carry. This is intrinsic to the Internet's success in developing so rapidly as a global infrastructure, as the administrative burden for a new network to become part of the Internet is relatively low. Just like the analogy of sharing maps in previously unknown territories, the Internet largely works because all networks freely share routing information with one another and the default assumption is always that this information is accurate.

The problem of inaccurate routing information is, in part, a legacy of the trusting nature of the relatively small research community which originally developed the Border Gateway Protocol. There are risks and uncertainties related to routing information on the Internet, and we argue that trust is an important way to alleviate such concerns without formal governance. In our examples, uncertainty deals with the reliability of information flows (e.g. will the information be accessible to others?) The risk is what is actually at stake in the transfer of information across the Internet. Risk is equivalent to the importance of information that would not be accessible or properly transferred to the intended target in the event of a failure.

Using a series of qualitative interviews and historical research, we argue that social relationships and the maintenance of trust are essential for safeguarding the stability of the Internet. We show how social relationships structured as exchanges amongst network administrators, along with a social structure following the structure of the Internet itself contribute to the maintenance of order despite the lack of formal governance. Furthermore, we discuss how the specific form of technology and associated technological practices contribute to the spread of trust transitively across the Internet, in turn driving the formation and maintenance of community and reputation to maintain stability. This is a story of entwined social and technical systems which together contribute to the maintenance of order in the absence of centralized authority.

## Data and Methodology

The research presented here is based on a variety of different materials, including vetted historical accounts of the creation of the Internet and interviews with network administrators. 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) history of the development of the Internet, as well as several other accounts of significant events in the history of the Internet.

From an ethnographic standpoint, this project poses an interesting challenge since the primary subject is a global network and there is no single, defined field site where research on the Internet infrastructure can be conducted. The distinctive elements of the phenomenon under study – its decentralized, placeless character, and the materiality of the technology – make it a potentially difficult subject. Just as the postal system is not best described by the vehicles and roads that carry the mail, it is the behavior, attitudes and beliefs of the people who operate and manage such complex systems that are of principal interest in this research.

The qualitative portion of our analysis is drawn from 23 interviews which were conducted in 2008-2010 to gain a first-person understanding of the world of Internet network administration. Interviewees in this study were primarily network administrators, but also include 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 the Internet's immediate ancestor from the late 1980's, the NSFNET, to others who began their careers as little as five years ago. Interviews were conducted by telephone, in person at the interviewees' offices, at the 74th IETF meeting, and at the 49th NANOG meeting, both of which were held in San Francisco, California. Interviews lasted from 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 snowball sampling was adopted, where current interviewees provided introductions to acquaintances as possible interview candidates.

## A Brief History of the Internet and Key Technologies

In order to best describe the relationship between Internet network administration and information routing, we begin with a brief history of the key technologies that are principal to our analysis. The technologies of the Internet have their origin in the ARPANET, a network funded by the Advanced Research Projects Agency (ARPA) of the United States' Department of Defense. The ARPANET was initially funded to provide access to ARPA-funded computing installations from remote academic sites. These technologies were later extended to provide "inter-networking" facilities connecting the ARPANET and the ARPA packet radio network. Commenting on the development of the Internet protocols in this period, Clark (1988) notes that:

> The components of the Internet were networks, which were to be interconnected to provide some larger service... At the time it was assumed that there would be other sorts of networks to interconnect... networks represent administrative boundaries of control, and it was an ambition of this project to come to grips with the problem of integrating a number of separately administrated entities into a common utility.

The ARPANET was succeeded by the NSFNET, a more modern network funded by the NSF intended to facilitate communications amongst research institutions across the United States, and to serve as a test bed for evolving Internet technologies. The Internet was born in 1994 when the NSFNET infrastructure

was privatized, and numerous independent networks interconnected to form a seamless whole (Abbate 1999).

The Internet that we know and use today is, therefore, an instance of an “internet”, a much larger network of various interconnected networks. Each of these networks optimizes data flows within itself across various technical parameters, such as overall bandwidth utilization, 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. 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.

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 neighboring Tier 2. Stub networks are those that exist at the periphery of the Internet, buying transit from Tier 1 or Tier 2 networks.

## *What is the Border Gateway Protocol (BGP) and why does it matter?*

The Border Gateway Protocol (BGP) provides the critical mechanisms to enable these disparate networks to interconnect with one another. A series of BGP specifications were standardized in the NSFNET era, from version 1 in 1989 to version 4 in 1995 (Rekhter and Li 1995), which is currently in use on the Internet. BGP version 1 was standardized 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).

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)[1], which in turn is granted blocks of ASNs for allocation by the Internet Assigned Numbers Authority (IANA)[2]. For instance, the autonomous system operated by the University of California, Berkeley is AS25 (Autonomous System number 25). There are currently over 35,000 autonomous systems active on the Internet (Huston 2010).

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.

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 autonomous system 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 Problem of Trust

We argued at the beginning of this paper that trust is inherent to the development and maintenance of components in the Internet infrastructure such as BGP. To elucidate this point, we must take a temporary detour from the technical descriptions and describe precisely what we mean by the term 'trust' in the context of the network administrators who maintain the Internet. For example, we do not mean trust in the technology itself (e.g., will the technologies work properly, what would make them fail, etc). Such discussions are a part of the wider literature on trust and technology (see Nissenbaum 2001, 2004; Marsh and Dibben 2003). We are less interested in trust in the technology as we are in trust among those who maintain the technology.

Social interaction in different situations entails varying degrees of risk and uncertainty. When two individuals interact for the first time and have something real to lose (e.g., a financial transaction), each party makes assessments, implicitly or explicitly, about the trustworthiness of the other. Risk is what is at stake in a given interaction (such as money or time), while uncertainty refers to the ambiguity of an outcome (Cook et al., 2005). Both risk and uncertainty influence one's assessment of another's trustworthiness in a given context. A highly uncertain interaction may not require much trust if the investments are very low, yet a near certain outcome could require enormous trust if what is at stake is significant. In the case of network administration, individual network operators might build and even

1 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.

2 http://www.iana.org/

maintain trust among those who operate adjacent networks. These administrators share many of the same risks and uncertainties regarding proper routing information through their respective networks and the transmission of inaccurate address advertisements. An interpersonal, human relationship is certainly not a given—but the key issue is that there is an incentive to build and maintain trusting relationships between network administrators when they rely on each other to identify and respond to risks when and if they occur.

An important insight from the vast literature on trust is that trust and trustworthiness are different but related concepts. Trust is an expectation of favorable reciprocity from others in situations that are uncertain or risky (Brann and Foddy 1987). On the other hand, trustworthiness is a characteristic of one who is trusted (Hardin, 2002). Sometimes individuals accept a risk with little or no information about those with whom they will interact. These one-time interactions tend to describe pure risk-taking rather than trust (Hardin, 1993; Hardin, 2002). Yet, there are situations where individuals make an evaluation and decide whether or not to take a risk with one or more individuals. Various factors affect our judgments of trustworthiness, including the nature of the situation, the mode of interaction, and perceptions about another's intentions and motivations (Cook et al., 2009). This initial action is an evaluation of trustworthiness. If the relationship continues with mutual reinforcement of risk-taking, the parties might be considered trusting of one another (Hardin 2002, Cook et al. 2005). To return to our earlier example of network administrators who are dependent on one another for the proper functioning of their respective systems, an assessment of trustworthiness might be an evaluation of one's competence or motivation to act in an expected way. These expectations might include proper maintenance and configuration of one's own network, as well as prompt response to problems or other risks. However, the evaluation of trustworthiness is simply a first step—only after ongoing interactions (e.g., fulfilled expectations about the behavior of one's peer network) might the two administrators build trust through experience.

The positive outcomes that we associate with the concept of trust can act as a type of social glue, alleviating concerns of risk and uncertainty in interpersonal interaction (Govier 1997). When individuals follow through on inter-dependent roles and obligations, one's prior actions act like a signal of future behavior. This creates what Axelrod (1984) calls the *shadow of the future*. One's current and past experiences with another person casts a figurative shadow of expectations over all future interactions with the same person. The ongoing interactional component of social relationships is what makes interpersonal trust possible, as trust cannot exist in one-time interactions (Hardin 2002; Cook et al. 2005; Cook, Hardin and Levi 2005). One can guess or approximate another's intention (assessment of trustworthiness), and even take a chance with that person (act of risk-taking). However, true interpersonal trust between individuals results from a process of mutual risk-taking over time (Cook et al. 2005). Through these behavioral experiences, individuals learn about one another's intentions and perhaps even their attitudes and beliefs. At minimum, trust depends on, "a fairly well defined interest at stake in the continuation of the relationship" (Hardin 2002: 3).

Just because individuals act in a cooperative manner, it does not necessarily mean that they trust one another. For example, when there are organizational or institutional mechanisms in place to shield individuals from betrayal, people tend to rely primarily on these third-party protections *instead* of building mutual trust (Cook, Hardin, and Levi, 2005). Thus, assurance structures that are designed to facilitate trust in uncertain environments can undermine the need for trust in the first place.

For example, consider an ongoing negotiation between network administrator A and adjacent network administrator B that is completely ensured by a 3rd party. The interaction may lead to cooperative, mutually beneficial outcomes for A and B. However, there is arguably no real trust relationship between A and B because the interaction is guaranteed by the 3rd party. More precisely, persons A and B need only trust the 3rd party, not one another. This case is entirely plausible if, for example, the two network administrators simply monitor their own systems but entirely rely on the technologies to maintain and fix themselves. In such a case the technologies would act as a 3rd party assurance system. Once the 3rd party

disappears, substantially more uncertainty will exist between A and B. Trust can emerge, when institutional mechanisms collapse, but history indicates that this is not a certainty (Cook, Rice and Gerbasi 2004). Individuals who become dependent on a 3rd party do not easily trade their trust in that party for trust in one another.

## *Building Trust into BGP*

In addition to the trust that can emerge between individuals, the assumption of trust can also be built into the technologies that route and transfer our information. As our earlier account of BGP illustrates, inter-domain routing functions something like a game of "Whisper Down the Lane", where players pass a phrase or sentence to one another to see if the original statement makes it intact to the final player. On the Internet, each autonomous system must trust that its peers are being truthful in the claims they make about routes which they can carry in their BGP advertisements. There is no mechanism for an autonomous system 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 'problem' of trust is designed into BGP itself, as the protocol offers no support for validating the claims that peers make in their advertisements. This issue is magnified many times over by the fact that peers often act simply as relays for advertisements from their peers. More often than not, 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. This, in effect, allows incidents such as the hijack of YouTube's IP address space by the Pakistan Internet Exchange.

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 solely the trust expressed between peers; it is also the transitive trust that establishes itself across the web of interconnected networks that make up the Internet. In human to human relationships, trust is not necessarily transitive. Specifically, trust cannot be truly transitive between people because, "its effective content is not derived from the formal characterization of its intrinsic properties" (Barbalet 2005: 11). But this is precisely the kind of trust that is built into the assumptions of BGP: to trust a peer network is to trust that network's peers, and so on. Trust relationships follow the structure of the network, and are in fact carried by this structure in the transitive trust mechanisms enforced by BGP. Trust in this instance is a quantity that is at once social (between administrators who manage separate, interconnected networks) and also engineered into one of the core protocols of the Internet (BGP) that allows networks to route information through and between one another. In essence, the particular form of the technology *forces* trust in the core systems of the Internet. As we will continue

to argue below, the assumption of trust embodied in the technologies that link networks together also contribute to a sense of community amongst those working in these core networks.

## The Social Origins of Inter-Domain Routing

The development of Internet technologies, and of the early Internet infrastructure itself, began as a research effort in the form of the ARPANET, and later the NSFNET. The ARPANET was developed in the 1970s, with funding from the Advanced Research Projects Agency (ARPA) of the United State's Department of Defense. 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. One of our interviewees 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. This is not to say that the IETF is an idealized meritocracy; its basis lies in thick relationships amongst a core of recurrent contributors. During an observed IETF meeting, for instance, in the midst of heated technical arguments at a working group, one of those speaking commented laughingly to the room at large that all those involved were good friends, and had known one another for a long time.

There is a high level of generalized trust in the Internet community, which provides the space for newcomers to contribute. However, this community (like many others) is also patterned with thick relationships formed over time, in which interpersonal trust is a critical element. It is important to note that this is the context in which BGP, and other Internet technologies, were developed.

The combination of a high level of generalized trust, and interpersonal trusting relationships resulted in closure on a form of BGP that was also “trusting”: mechanisms to validate route advertisements were never built into the protocol. In our interviews, a common refrain has been that the trust engendered by the close ties amongst the Internet community was 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 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.

The hierarchical structure of the NSFNET allowed the NSFNET backbone administrators control over the routes that NSFNET would carry, ensuring that erroneous 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 said:

> 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 that a protocol essentially intended for the political purpose of separating network administrations should be secured, that it should not be quite so trusting by default.

In fact, the trust amongst the members of this tightly knit community played a central role in the early Internet, a theme which many of our interviewees commented on:

> 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 the research and operational elements of the Internet community 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, as of 2010, remains with us.

## *From NSFNET to Internet*

Since NSFNET only allowed non-commercial traffic, a number of regional networks arose to support commercial demand. Recognizing 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 privatization 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.

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, perhaps, 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 recognized, 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)[3], 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 organization 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 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 interviewee 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

3 See http://www.ra.net

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 our 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.

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[4], of this delegation. During an interview, 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' 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 organization for checking such requests before modifying BGP configurations to route the requested address space:

> ... we're usually pretty diligent about checking the SWIP[5] 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 interviews, illustrating how these challenges have changed the administrative practices around managing IP address blocks:

---

4 IANA – Internet Assigned Numbers Authority, RIR – Regional Internet Registry.

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

> 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, sent in a letter to ARIN[6] 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 anymore 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 standardizing 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 attitudes towards 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, and is regarded as a best practice in the network administration community (Ferguson and Senie 2000). 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, an issue to which we return in the next section.

## Creating Order on the Internet

Words like "community", "family" and "trust" came up quote frequently in our 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 did not hold true anymore: 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 stated:

> 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.

---

6 The Regional Internet Registry for North America.

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 our 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 interviewee 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[7] before that, both of those were done by pointy-haired bosses, that are that type of people that were looking at the bottom line, and didn't really realize what havoc they were going to wreak for their *own business*. They weren't just causing other people problems, they were causing their own people problems.

As these quotes suggest, our interviews indicate a cohesive social world that extends across the networks – commercial and non-commercial – that make up the Internet. We began our research in search of this community, expecting it to be relatively undifferentiated, and mediated through the NANOG email list and meetings. We 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 who we interviewed had no sense of being part of the larger Internet community, since they connected to the Internet through a single provider network; their relatively small size does not justify having redundant connections through multiple providers. As a result, administrators at small ISPs are more involved in local professional groups of systems administrators, rather than NANOG (though they may well listen in on the NANOG email list). 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. One such administrator commented that “as far as receiving routes, I take whatever my upstreams send me, and pretty much take it on blind faith.” To trust is to take a risk in making one choice over another; relationships which form in the absence of choice can only be said to involve confidence, rather than trust (Luhmann 2000). Since administrators at small ISPs lack choice in

---

7 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.

their BGP configurations, they have confidence, rather than trust, in their provider networks. 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[8] 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 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. 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, we 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[9]. 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 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 judgments 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.

---

8 Bay Area Regional Research Network, one of the NSFNET's regional networks.

9 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.

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:

> 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 [sic] 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 our interviewees commented that attendees generally tend to be personnel from Tier 1, Tier 2 and academic networks. When an administrator at a Tier 1 network was asked 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, described 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. In talking about practices of filtering, for instance, a Tier 1 network administrator said 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.”

One of our 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.

In addition, reputation is based on the relative position of a network in the Internet. When an administrator at a Tier 1 network was asked 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 our interviewees operates a network that is, for all intents and purposes, a Tier 1 system. 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.

## Trust and the Governance of the Internet

Governance is often treated as a problem to be addressed with the Hobbes's (1651) solution of investing in strong institutions (the so-called "Leviathan") to maintain order in a given setting. While this is certainly an important element of any discussion of governance, it is insufficient in the face of a system based on an assumption of decentralization such as the Internet where disparate political interests, economic motivations and other social forces rub up against one another as peers and competitors. Granovetter (1985) observes that Hobbes provides an account of governance that is too simplistic, where anarchy is undersocialized (e.g., anonymous, lacking human social relationships) and the Leviathan is oversocialized (a maximally connected, controlled, social environment). As Hardin (2002) argues, "Hobbes may have exaggerated the extent to which powerful institutional sanctions are required for grounding trust and promises, but he was not radically mistaken" (187). To apply such thought to the problem of governance and the Internet infrastructure, ICANN and the RIRs may set policy for the distribution of resources, but we believe that it would be unwise to depend on a singular global institution as the best and only solution to ensure order in the management of these resources. This is especially true in the critical arena of the inter-domain routing system which keeps track of which networks maintain which IP address blocks, and the intermediary networks which may be used to route traffic to them.

Russell Hardin (1998) makes a critical point that trust in a governing body is fundamentally not the same as trust between individuals. He argues that those who are governed need not trust the governing institution, and "all that is needed …is for citizens not actively to distrust it" (Hardin 1998: 11). Individuals do not or cannot actually trust those who govern because they have no true interactions with the people, only with their policies (Hardin 1998). When we relate this point to the Internet infrastructure, individuals who depend and rely on the Internet for their daily business do not necessarily need to trust those who maintain it:

> A claim to trust government is typically implausible if it is supposed to be analogous to a claim to trust another person…The difficulty with 'trusting government' is that the knowledge demanded by any of [the] conceptions of trust is simply unavailable to ordinary citizens. (Hardin 2002: 151).

The greater populace simply does not interact directly with those who govern and maintain order. On the other hand, those who govern individual resources (network administrators in our case) directly interact with one another to uphold the system. Much like medieval vassals with no king, individual administrators govern their respective networks and trust others with whom they directly interact to do the same. Rather than creating a specific type of order, it might be more accurate to say that the Internet is maintained by administrators who actively attempt to avoid disorder. Furthermore, because they are equally dependent on one another to route information properly, they have a shared interest in finding and solving problems when and if they occur.

We have argued that the technology of inter-domain routing, BGP, embeds a transitive form of trust as a consequence of the social context of its development. This transitive trust in the technology, in

combination with the social norms that encouraged the development of trust in the earlier NSFNET period, contribute to the sense of community amongst network administrators that is a core part of the Internet today. In fact, this type of generalized trust across network administrators in similar roles may be an example of what Farrell (2009) describes as class-based trust. Individuals can build trust with one another even when they do not have personal relationships if they have a broader form of knowledge about how actors in their same class or set of roles should behave (Farrell 2009: 133). An alternative but related network view is that dyadic trusting relationships between pairs of group members (such as adjacent network administrators) motivate cooperative action within the larger group. What we observe as generalized trust among those in a certain class or type of role is really just distributed cooperation produced from the many dyadic trust relationships within the group. Thus, "…trust relationships can enable us to cooperate beyond the dyadic level to some extent even though such relationships and their effects must run out for interactions within very larger groups" (Hardin 2002: 182).

Gellner incorporates trust in a compelling alternative to the Hobbesian model of social order, drawing from 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 wrote in relation to 14th century Arab society, the inter-domain routing system may well fall within the range of real human conditions that Gellner suggests would fit this model. As we 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 quote from the NANOG email list makes this point clear:

> 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. Its part of what makes us an *autonomous* system, sovereign of our own little kingdom.[10]

In this setting, there can be no substitute for interpersonal trust, and the attendant sense of community which allows for a decentralized mode of governance over an essentially distributed system. In the context of the Internet's inter-domain routing system, trust is coextensive with the maintenance of socio-technical order. When strong institutional forces are present, this expectation of trust among similar others might seem inconsequential. As we have shown through the perspective of network administrators, however, it is precisely due to the *absence* of strong institutional oversight and regulation that interdependent trust relationships and generalized cooperation were able to emerge over time.

## IMPLICATIONS AND CONCLUSION

Consider a counterfactual Internet, one dependent on a central authority for the management of routing information (much like the NSFNET). Where ICANN, perhaps, took on the mantle of a trusted third party which would warrant all BGP route advertisements. What would such an Internet look like? Following the recommendations being developed at the IETF's Secure Inter-Domain Routing (SIDR) Working Group, the central authority would maintain the root key servers for the public key infrastructure required to offer these warrants in a distributed manner. This would create a form of trustworthy computing, where

10 From http://www.merit.edu/mail.archives/nanog/msg17353.html

trust would devolve from interpersonal relations to a singular institutional context. Such an architecture may result in a more secure or reliable Internet, since it would become virtually impossible to issue false route advertisements, through misconfiguration or malice.

The choice between assurance structures through technological solutions or trust relationships without structural assurances is fundamentally a tradeoff between different but related problems. Without assurance structures individuals can, but not necessarily will, build trust as a solution to the problem of risk and uncertainty. With assurance structures in place, individuals behave as if they trust one another—but this behavior is unstable and should not be mistaken for genuine interpersonal trust. For example, in cross-societal studies of trust relationships, researchers have noted that social relationships in Japan depend on formal commitment mechanisms (assurance structures) as a way of avoiding the risk associated with exploitation (Yamagishi 1998; Yamagishi, Cook and Watabe 1998). Indeed, in Hofstede's (1980) cross-cultural examination of values, the Japanese were shown to be more likely than Americans and many other societies to favor situations with lower social uncertainty. This problem of assurance versus trust has been demonstrated in experimental research that shows that the Japanese are less willing to engage in cooperative behavior in the absence of an assurance system compared to Americans (Yamagishi 1988, Yamagishi, Gillmore and Cook1988; Cook et al. 2005). The key point from this line of research is not simply that cross-societal differences exist. Rather, the crucial implication is that those who become accustomed to assurance structures come to rely and depend on them at the expense of trust. When and if these assurances fail, cooperation and trust fail as well.

A second problem with the hypothetical central authority for the Internet is that it would become a central point of political contention from an institutional perspective, and a central point of failure from the perspective of technological infrastructure. An administrator at a Tier 1 network related his concerns to us, "if you break the root key server, now you break all routing in my network, which is a much more difficult problem for me to troubleshoot [...] and there's always the fear that one particular government might take control over what's allowed to be published in the key server." A strong central authority makes for a potentially brittle system of governance where technological failures may cause global routing problems. In addition, adverse political interests may be left with no option but to create their own parallel institutional and technological structures, effectively partitioning the Internet. The recent leak of routing information from China, "importing censorship" by redirecting DNS requests from around the world to Chinese DNS servers[11] provides a soft example of how such a system might play out in practice.

The opposite extreme would be to function with no central authority, and no technological mechanisms for warranting route advertisements. This is, in many ways, the Internet as it is today, which we have described in our account. The lack of warranting mechanisms in BGP could be viewed as an elegant solution to the difficult problem of allowing for easy attachment of new networks to the Internet, and for relatively simple ways to change routing announcements. Indeed, it could be argued that the Internet would not have enjoyed the extent of growth that it has had thus far if it were not for this particular trusting technological form. The social world which we have described in our account is an essential element of this system, providing for local autonomy while at the same time actively avoiding global disorder. It is a system that works, in part, because it is small enough to operate on norms of cooperativeness across the collection of network administrators. Furthermore, the dyadic connections between adjacent network administrators allow trust to develop and persist over time. These are mutually reinforcing phenomena, as the network administrators in our study indicate through their personal accounts of the larger cooperative environment and the value of trust among one's peers. This is precisely the distinction that Hardin (2002) makes when he states that, "When we have the thick relationships of a small, close community, we may find that our interactions are governed by norms of cooperativeness that are collective rather than dyadic" (183). However, Hardin also notes that there is a type of compromise between groups that are built on norms of cooperativeness:

11 See http://www.renesys.com/blog/2010/03/fouling-the-global-nest.shtml

> An ideal compromise world might be one in which we have relatively thick relations with some core group or groups and then far-flung networks of relationships with respect to many and varied particular things. With quasi thick relations, we might better be able to translate the trustworthiness that we develop within a relatively close group or neighborhood to other contexts that are more limited in their range of interactions. (Hardin 2002: 185).

The Internet today, and for the foreseeable future, is defined by tussles amongst adverse interests: nation states, corporations, public interest groups and more. Clark et al (2002) suggest that the technologies of the Internet should be built in a manner that allows these tussles to play out as they will, with an explicit attempt to avoid encoding political inclinations of one sort or another into technological form. Similarly, we argue that the *institutions* of the Internet should be designed in a manner which allows for tussle, for different local implementations of the Internet to interoperate and exist alongside one another. Such a system must rely on a hybrid form of governance, which balances centralized authorities with decentralized local agency. In putting this position forward, we are not arguing against enhancing Internet technologies to create secure and reliable computing systems. Rather, we are arguing for a combined focus on technological and institutional design, with a consciousness of social forms and everyday practice in network administration. Our goal should be to create trustworthy operating *environments* (rather than solely trustworthy/reliable/secure computing), in which relations of trust amongst disparate actors may be encouraged, cultivated, and esteemed.

## Acknowledgements

We would like to extend our thanks to all of our interviewees, who generously gave their time to this project. Our thanks also go to the attendees of the 74th IETF and the 49th NANOG meetings for providing an intellectually stimulating atmosphere for our research. Finally, we salute the network administrators of the world, whose invisible work keeps the Internet running.

## References

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

Axelrod, Robert. 1984. *The Evolution of Cooperation*. New York: Basic Books.

Barbalet, Jack. 2005. Trust and Uncertainty: The Emotional Basis of Rationality. Taking Stock of Trust Conference, London School of Economics. December, 2005.

Brann, P. and Foddy, M. 1987. Trust and the consumption of a deteriorating common resource. *Journal of Conflict Resolution*, 31, 615-630.

Carr, S., S. Crocker and V. Cerf. 1970. Host-Host Communication Protocol in the ARPA Network. In: Proc. AFIPS SJCC Vol. 36.

Clark, D. 1988. The design philosophy of the DARPA internet protocols. In *Symposium proceedings on Communications architectures and protocols*, 106-114. Stanford, California, United States: ACM. doi:10.1145/52324.52336. http://portal.acm.org/citation.cfm?id=52336.

Clark, David D., John Wroclawski, Karen R. Sollins, and Robert Braden. 2002. Tussle in cyberspace: defining tomorrow's internet. In *Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications*, 347-356. Pittsburgh, Pennsylvania, USA: ACM. doi:10.1145/633025.633059. http://portal.acm.org/citation.cfm?id=633025.633059.

Cook, Karen S., Chris Snijders, Vincent Buskins, Coye Cheshire. 2009. *eTrust: Forming Relationships in the Online World*. New York, New York: Russell Sage Foundation.

Cook, Karen S., Eric R. W. Rice, and Alexandra Gerbasi. 2004. The Emergence of Trust Networks under Uncertainty: The Case of Transitional Economies--Insights from Social Psychological Research. In *Problems of Post Socialist Transition: Creating Social Trust*, edited by Susan Rose-Ackerman, Bo Rothstein, and Janos Kornai, pp. 193-212. New York, NY: Palgrave Macmillan.

Cook, Karen S., Russell Hardin and Margaret Levi. 2004. *Cooperation without Trust?* New York, New York: Russell Sage Foundation.

Cook, Karen S., Toshio Yamagishi, Coye Cheshire, Robin Cooper, M. Matsuda, and R. Mashima. 2005. Trust Building via Risk Taking: A Cross-Societal Experiment. *Social Psychology Quarterly*. Vol. 68, Number 2, pp. 121-142.

Farrell, Henry. 2009. Constructing Mid-Range Theories of Trust: The Role of Institutions. In *Whom Can We Trust? How Groups, Networks, and Institutions Make Trust Possible*. Karen Cook, Russell Hardin and Margaret Levi, Eds. New York, New York: Russell Sage Foundation.

Ferguson, P., and D. Senie. 2000. BCP38: Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing. May. http://www.faqs.org/rfcs/bcp/bcp38.html.

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

Govier, Trudy. 1997. *Social Trust and Human Communities*. Montreal, Canada: McGill-Queen's Press.

Granovetter, Mark. 1985. Economic Action and Social Structure: The Problem of Embeddedness. *The American Journal of Sociology* 91, no. 3 (November): 481-510.

Hardin, Russell. 1998. Trust in government. In *Trust and Governance*. Valerie Braithwaite and Margaret Levi (eds). New York: Russell Sage Foundation, pp. 9–27.

_____. 1993. The Street-Level Epistemology of Trust. *Politics and Society* 21, 505–529.

_____. 2002. *Trust and Trustworthiness*. Russell Sage Foundation: New York City, New York.

Hobbes, Thomas. 1968 [1651]. *Leviathan*. Harmondsworth, England: Penguin.

Hofstede, Geert H. 1980. *Culture's Consequences, International Differences in Work-Related Values*. Beverly Hills: Sage.

Huston, Geoff. 2010. CIDR Report. August 15. http://www.cidr-report.org/as2.0/.

Ju, Qing, Beichuan Zhang, and Varun Khare. 2010. Large Route Leak Detection presented at the NANOG49, June, San Francisco. http://nanog.org/meetings/nanog49/abstracts.php?pt=MTYwMiZuYW5vZzQ5&nm=nanog49.

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

Nissenbaum, H. 2001. Securing Trust Online: Wisdom Or Oxymoron. *Boston University Law Review*, 81, 635.

_____. 2004. Will security enhance trust online, or supplant it? In R. Kramer & K. S. Cook (Eds.), *Trust and Distrust Within Organizations: Emerging Perspectives, Enduring Questions,* pp. 155–188. New York, NY: Russell Sage Publications.

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.

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

Powell, Walter W. 1990. Neither Market nor Hierarchy: Network forms of organization. *Research in Organizational Behavior* 12: 295-336.

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

Weber, Max. 1958. *From Max Weber: Essays in Sociology*. Browning of Pages. Oxford University Press (Galaxy imprint).

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

Yamagishi, Toshio. 1988. The provision of a sanctioning system in the United States and Japan. *Social Psychology Quarterly* 51: 32-42.

Yamagishi, Toshio, Karen S. Cook and Motoki Watabe. 1998. "Uncertainty, trust and commitment formation in the United States and Japan." *American Journal of Sociology* 104: 165-194.

Yamagishi, Toshio, Mary R. Gillmore and Karen S. Cook . 1988. Network connections and the distribution of power in exchange networks. *American Journal of Sociology* 93, pp. 833–851.