

Mathew, Ashwin J.

Article

The myth of the decentralised internet

Internet Policy Review

Provided in Cooperation with:

Alexander von Humboldt Institute for Internet and Society (HIIG), Berlin

Suggested Citation: Mathew, Ashwin J. (2016) : The myth of the decentralised internet, Internet Policy Review, ISSN 2197-6775, Alexander von Humboldt Institute for Internet and Society, Berlin, Vol. 5, Iss. 3, pp. 1-16,
<https://doi.org/10.14763/2016.3.425>

This Version is available at:

<https://hdl.handle.net/10419/214020>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by/3.0/de/legalcode>



The myth of the decentralised internet

Ashwin J. Mathew

School of Information, University of California, Berkeley, United States, ashwin@sanmathi.org

Published on 30 Sep 2016 | DOI: 10.14763/2016.3.425

Abstract: In popular culture, and in policy discussions, the internet is often conceived of as a decentralised technology, which cannot be controlled. Drawing from research into internet infrastructure, focusing on the Border Gateway Protocol, I show that the internet has never been, and never can be, decentralised. I argue that the internet is better viewed as being distributed, both in terms of technologies and governance arrangements. The shift in perspective, from decentralised to distributed, is essential to understand the past and present internet, and to imagine possible future internets which preserve and support the public good.

Keywords: Decentralisation, Governance, Protocols, Routing

Article information

Received: 25 Apr 2016 **Reviewed:** 17 Jun 2016 **Published:** 30 Sep 2016

Licence: Creative Commons Attribution 3.0 Germany

Funding: This work was funded in part by the TRUST Center, under NSF award CCF-0424422. The writing of this paper was supported by a grant from the UC Berkeley Center for Long Term Cybersecurity, and by a fellowship at the Slow Science Institute.

Competing interests: The author has declared that no competing interests exist that have influenced the text.

URL: <http://policyreview.info/articles/analysis/myth-decentralised-internet>

Citation: Mathew, A. J. (2016). The myth of the decentralised internet. *Internet Policy Review*, 5(3). DOI: 10.14763/2016.3.425

This paper is part of Doing internet governance, a special issue of Internet Policy Review guest-edited by Dmitry Epstein, Christian Katzenbach, and Francesca Musiani.

INTRODUCTION

From its earliest days, the internet has been imagined as a space which holds the potential for a freer and more democratic form of social organisation, eliminating the inequalities, regulations and controls of societies run by governments and corporations. This social imaginary of the internet relies upon a key assumption: that the internet cannot be controlled by any centralised or territorial authority. Concerns that the internet is being subjected to centralised control and regulation in the present invoke a past in which it was free from corporate and government control, as "a world that is both everywhere and nowhere" (Barlow, 1996), and a future in which it may be free again, constructed through technologies which are inherently "decentralised,

radically democratic” (Granick, 2015). Following this line of analysis, the most enduring responses to attempts to control the internet involve designing new technologies which are ever more decentralised, capable of evading control.

The early internet did appear decentralised to its users, as they were able to create and use new services without needing to seek permission from any third party. I share the concerns that the everyday experience of the internet of today is increasingly controlled, through concentrations of power in the ownership of media and platforms (Benkler, 2016; Lessig, 2001). However, as I will show, the experiences of apparent decentralisation and control are both constructed over an underlying infrastructure which was never decentralised, nor designed with decentralisation as a goal.

I argue that the internet is better conceived of as a distributed system – rather than a decentralised system – with varied centres and concentrations of power in its construction. This is not simply a matter of semantics. A distributed imaginary of the internet calls to attention institutions and practices, just as much as it does technology. In this analysis, responses to concentrations of power cannot focus on technology as a mechanism for eliminating centralised control. Instead, socially desirable outcomes (such as freedoms of speech and association) must be addressed through appropriate combinations of political and technological interventions.

To make my argument, I examine the origins and evolution of a key technology of the internet: the Border Gateway Protocol (BGP). BGP is the technology which enables the interconnection of separate networks to form the global internet. After all, the internet is not one network, but an interconnected system of tens of thousands of individual networks.¹ BGP enables the routing of data between these individual networks. BGP is one of the most critical components of internet infrastructure: in the absence of the interconnections between networks, there would be no internet. Since the routing of data on the internet is managed through BGP, the analysis of the internet as decentralised or distributed is in many ways defined by the characteristics of BGP, making this protocol an essential object in the study of internet governance.

I study BGP through three related technological features which are fundamental to a decentralised conception of the internet. First, packet switching, the mechanism through which data is broken up into individual packets before transmission. Second, routing, the means through which routes are discovered and selected for the transmission of packets within and between networks. Third, topology, the structure of the system of interconnections between the networks which make up the internet. These three features must work in concert if a decentralised internet is to be possible: communications are broken up into packets, which may take different paths – chosen using a decentralised routing protocol – through the topology of the internet. By examining these three features, I will show that decentralisation was not a design goal, nor the actual outcome, in the creation and subsequent operation of BGP, and by extension, of the internet; and I will illustrate why it is important to understand the internet as a distributed – rather than decentralised – system.

I begin the paper with a brief review of debates in the governance of technology, focusing on problems of control in the development and operation of infrastructure. In the following sections, I trace the evolution of the relationships between technological form, control and topology which were required to govern internet routing. First, I examine the internet's earliest incarnation, the ARPANET, in which routing was administered by a single organisation through centralised control. Next, I follow the transition of routing in the ARPANET's successor, the NSFNET, which encompassed multiple independently administered networks, with routing managed through a hierarchical network topology, and corresponding hierarchical control.

Finally, I trace the operation of routing in the modern internet, in which topology takes the form of a complex graph, administered through coordination between varying centres of power.

The historical approach that I adopt serves two purposes. First, it provides a comparative case study of changing forms of governance across different periods in the history of the internet, through distinct articulations of technological form, control and topology. My analysis and method, I hope, will offer generalizable lessons for the study and design of internet technologies and governance. Second, it traces the values and assumptions underlying the evolution of BGP as a socio-technical system: values and assumptions which in many ways remain embedded in the modern form of BGP. To make sense of the possibilities that the internet offers, it is essential to understand the mechanisms through which its infrastructure – of which BGP is a critical component – is governed.

SECTION 1: GOVERNING TECHNOLOGY

The problem underlying both centralised and decentralised visions of the internet is one of governing technology: what mechanisms are required to stabilise and order the global technological infrastructure of the internet? The debate is over the form that these mechanisms take, in promoting centralised control, or decentralised coordination; and over the social values that users are able to realise in the resulting system.

While the subject of this debate – the internet – may be relatively new, the terms of the debate are not. Mumford (1964) argues that authoritarianism and democracy breed particular types of technologies, which encode the characteristics of the political system of their origin, providing for centralised/authoritarian or decentralised/democratic control. Winner (1980) presents technologies as a complex of features, with different capacities for flexibility or rigidity in relation to the patterns of power and authority which make use of them, and which are needed to govern them. Like Mumford, the cases he presents rely on the distinction between authoritarian and democratic politics, although with contingent connections between technological and political form.

There are sharp distinctions to the debate over the relationship between technology and governance, especially with regards to modern societies in which information is said to be a dominant factor. Beniger (1986) argues that centralised control, implemented through new bureaucratic administrative forms and the use of information technology, is essential to organising industrial production, telecommunications, transport and other large-scale elements underpinning modern "information" societies. In contrast, Galloway contends that the internet has created the conditions for the emergence of protocols, technologies of control which "cannot be centralised" (Galloway, 2006, p. 11).

In these perspectives, the governance of technology is constituted through the interactions between administrative structure and technological form. Governance in a centralised system functions through central bureaucratic administrative structures and technologies of control that govern from above. Governance in a decentralised system relies on relationships between individuals in flat social structures, enabled through technologies which support connection and coordination, without any central control. However, these positions are inadequate to explain the nature of the internet. For my analysis, I supplement administrative structure and technological form with a third analytical category of topology: the arrangement of relationships between entities in the system.

The object of my study, BGP, is part of the internet's infrastructure, invisible to internet users under normal conditions (Edwards, 2003; Star, 1999), but essential to the stable, reliable operation of the internet. I draw from, and extend, two dominant perspectives on infrastructure for this paper. The first focuses on the role of communities of practice, operating in situated contexts, in the use and design of infrastructures (Bowker & Star, 2000; Star, 1999; Star & Ruhleder, 1996). The second engages with the historical evolution of technology in large technical systems (Coutard, 1999; Hughes, 1983, 1987; Mayntz & Hughes, 1988), building on the stages of development established in Hughes (1983), from invention and development, to technology transfer, to system growth, to the acquisition of technological and political-economic momentum, to the rise of planned regional systems.

As valuable as these theories of infrastructure are, they fail to adequately develop a concept of the topology of infrastructure: the question of how the structure of the networks in which infrastructure is deployed (e.g., interconnected systems of canals, power distribution lines or roads) interact with the development of practices, standards and political economy of infrastructure.

Studies of internet governance provide a starting point for developing a topological approach to infrastructure. It is well established that internet governance is constructed through social networks and practices of ongoing coordination and collaboration, as much as through formal institutional structures (Hofmann, Katzenbach, & Gollatz, 2016; Mueller, 2010; van Eeten & Mueller, 2013). The aim of this strand of internet governance research is to understand the mechanisms through which distributed practices of coordination and collaboration constitute systems of governance.

My contribution to these theories of development and governance of technological infrastructures is to take topology as a central problem in the analysis of governance, to understand how coordination, collaboration, and power relationships function through topological positions and structures. I consider how the topological forms of internet infrastructure interact with the practices and social formations involved in operating internet infrastructure; and how these interactions structure the governance of internet infrastructure. I conceive of governance of infrastructure as premised upon situated, relational knowledge and practices (Haraway, 1988; Lave & Wenger, 1991), constructed through specific positions in the topology of infrastructure.

In my analysis, the power and authority required to engage in governance flow from topology (Allen, 2011; Harvey, 2012). Governance is not a global constant, but is constructed across varying related situated contexts in the topology of infrastructure. Situated contexts of governance are as much a matter of actual topological positions of actors within an infrastructure, as they are of the orbits of relations around actors, the circulation of practices involved in operating infrastructure, and the technological forms of the infrastructure.

SECTION 2: METHODS AND DATA

My research involved historical study into the origins and operation of BGP. I focused on the development of problems and technologies related to network interconnection, from the 1970s when these problems and technologies were initially defined, to the mid-1990s, when BGP matured as the key technology for network interconnection on the public internet. I analysed publicly available standards and best practice documents, and email lists from the internet

Engineering Task Force (IETF), which sets technical standards for the internet, focusing on material related to BGP and prior related technologies. These included documents from the IETF standards process, as well as Internet Engineering Notes (IENs), an earlier form of internet standards document. ² I studied historical and contemporary computer science papers related to BGP, and drew on first- and second-hand historical accounts of the development of the internet.

Over the course of a larger project – an ethnography of technical personnel involved in the operation of BGP – I conducted over 30 semi-structured interviews with individuals who were involved in operating and building the early internet. I draw from these interviews to provide first-hand accounts of the development of BGP and related technologies. I recruited interviewees through cold calls (and subsequent references) at meetings of the IETF, and of the North American Network Operators Group (NANOG), the principal professional organisation for the internet's technical personnel in the North American region.

Others have worked with similar material to study the manifestation of social values in the internet standards process (Braman, 2011) or debates in standards development around the transition from one internet protocol standard to another (DeNardis, 2009). In contrast, I am not concerned with the standards development process, so much as with the connection between specific features of a standard, the viable range of topological forms for related internet infrastructure, and the forms of governance required to operate this infrastructure. As a result, my focus in selecting interviewees and materials, and in analyzing this data, was as much on the rationale behind particular protocol design decisions, as it was on the practices of those involved in building and operating the infrastructure implementing a protocol.

SECTION 3: ROUTING PACKETS ACROSS THE INTERNET

The transmission of data across the internet takes place using a mechanism known as packet switching, implemented in a core technology called the Internet Protocol (IP). Older circuit switched networks required the entire circuit between two endpoints to be reserved for the duration of a communication (as for a phone call). In contrast, packet switched networks break up communications into individual packets of data, each of which can potentially take different paths to their eventual destination, at which they are reassembled. Claims that the internet is intrinsically decentralised often rest upon the capabilities of packet switching, through which packets are able to take alternate paths as required, naturally adapting to congestion and failures in internet infrastructure.

However, packet switching does not in itself support decentralisation. Two additional conditions are required: there must be a diversity of routes across the internet sufficient to avoid centralised control of routing, and there must be a sufficiently robust decentralised method for constructing maps of these routes.

To reach destinations on the internet, packets must traverse multiple independently administered networks, with routes determined by the routing protocols and policies applied to the routers (the network components handling routing of data) in these networks. Routing protocols define the standardised mechanisms through which routers learn about routes on a particular network, and specify methods through which administrative policies may be applied to the construction of routes. There is substantial complexity in the traversal of routes across the internet. The routers traversed may have different technical capabilities, and have different

behaviours, depending on how they are configured, their manufacturer, and their specific model. Different networks may use different routing protocols within their borders, running over network designs customised to their needs.

In order to interconnect networks across their borders – while allowing for variances within networks – a common routing protocol must be used. The Border Gateway Protocol (BGP) is the routing protocol which integrates diverse networks (termed "autonomous systems" by BGP) into the global whole of the internet. The resulting system of interconnections amongst networks is called the inter-domain routing system. The topology of the interconnected networks in the modern inter-domain routing system is a complex graph, consisting of over 55,000 individual networks as of this writing.³

If the internet is intrinsically decentralised, then it is in the inter-domain routing system that characteristics of decentralisation should be most visible. If the topology of the inter-domain routing system is not decentralised, then all packets within global or regional contexts will have to traverse a small set of networks, allowing these networks global or regional control over flows of data on the internet. If the routing protocol – BGP – for the inter-domain routing system is not decentralised, global or regional authorities will be able to control the routes that packets take across the internet.

The technological form of routing protocols, the topological organisation of networks, and related control mechanisms are key concepts for the analysis of the inter-domain routing system. In the following sections, I deploy these concepts to analyse the evolution of the inter-domain routing system, from the ARPANET to the present day.

SECTION 4: ADMINISTRATIVE CONTROL IN THE ARPANET

The internet has its beginnings in the ARPANET, a network funded by the Advanced Research Projects Agency (ARPA) of the US Department of Defense, although developed in an academic research setting. The precursor to the inter-domain routing system was a system of interconnections between the "gateway" computers which provided remote access to the networks of the organisations connected to the ARPANET. The first gateways for the ARPANET – called Interface Message Processors (IMP) – were installed to connect four sites in 1969. They were all constructed on the same template, with simplicity and reliability in mind, using identical hardware and software. Each IMP provided connectivity for a time-sharing host computer that was connected to it. Users logged in to their local host computer could connect to remote sites through the local IMP, which transmitted packets to their destination host computers through one or more intermediary IMPs (Abbate, 1999).

The initial implementation of routing in IMPs was "adaptive", seemingly decentralised. Each IMP made decisions about how to route packets based on the quality of links to neighboring IMPs. Changes in link quality were constantly measured by each IMP, and propagated around the network as input for independent routing decisions at each IMP (Heart, 1969). As the ARPANET evolved, the design for IMPs was generalised into a standardised specification for gateways, which used a routing protocol called the Gateway-to-Gateway Protocol (GGP) to determine connectivity, and compute routes between gateways (Hinden & Sheltzer, 1982; Strazisar, 1979; Strazisar & Perlman, 1978).

The development of adaptive routing on the ARPANET faced interlinked technical and administrative challenges. The ARPANET developers were faced with technical problems such as routing loops: conditions under which the failure of a link in one segment of the network led to packets being continually forwarded in a closed loop without being sent on to their destination (Strazisar & Perlman, 1978). New generations of IMP hardware were designed with improved capabilities, but had to be integrated to co-exist with older IMPs. It was not feasible to simply replace all the older IMPs, especially as the ARPANET grew to interconnect more sites, and became a utility for everyday use, even while it remained a platform for research and development of the nascent internet protocols. The design of the ARPANET routing protocols had to be pragmatic, to take into account the material and organisational impedances of a diverse network interconnecting multiple institutions, built over potentially unreliable physical telecommunications infrastructure.

Even though the GGP was designed to be adaptive, the diverse issues involved in operating the ARPANET called for centralised administration. Host computers remained under the control of their institutions, but the system of IMPs and their interconnections was centrally monitored and managed, under a contract from ARPA to the firm of Bolt, Beranek and Newman (BBN). The centralised management of this system became ever more important to the reliable operation of the ARPANET, as it grew and became more complex, accommodating multiple generations of IMP hardware and software which had to be able to inter-operate. As a specification for the design and operation of gateways noted: “For reasons of maintainability and operability, it is easiest to build such a system in an homogeneous fashion where all gateways are under a single authority and control, as is the practice in other network implementations.” (Hinden & Sheltzer, 1982, p. 3)

Even though the ARPANET was designed as a system in which gateways adapted to network conditions by taking independent routing decisions, the management of the ARPANET remained a centralised function. The apparent contradiction in routing protocol implementation – between decentralised decision-making at gateways, and centralised control of the system – was essential to the reliable operation of the ARPANET.

SECTION 5: TOPOLOGICAL CONTROL IN THE NSFNET

Similar contradictions were necessary for the reliable operation of the ARPANET's immediate successor, the NSFNET. The development of the NSFNET began in 1985, with funding from the US National Science Foundation. The NSFNET took a substantially different form than the ARPANET, with a three-tiered hierarchical network topology: networks at institutions within a geographical region were connected by regional networks, which in turn were connected by the single NSFNET backbone network across regions (Harris & Gerich, 1996).

A new routing protocol was devised for the NSFNET, the Exterior Gateway Protocol (EGP). It was so named to distinguish it from "interior gateway protocols", the routing protocols which operated within networks, behind gateways, to provide IP-based connectivity between computers in the same network. EGP termed these individual networks “autonomous systems”. Administrators for each autonomous system had absolute control to choose appropriate internal network topologies and interior gateway protocols to serve their institutional needs. However, connectivity to other autonomous systems on the NSFNET required the use of EGP, mediated through the regional networks, and the NSFNET backbone.

The separation of domains of governance in EGP was intended to support the independent evolution of individual autonomous systems. There was no need for centralised management and coordination of hardware and software upgrades to ensure interoperability, as was the case with the ARPANET. Ideas of autonomy were consciously designed into EGP to ensure that the NSFNET would remain flexible and open to continued development in a variety of dimensions, including hardware, software, and routing algorithms (Rosen, 1982).

Although networks had autonomy in internal matters, the NSFNET was still a hierarchical topological structure, operated with centralised control. Regional networks controlled routing among institutions within their regions. The NSFNET backbone controlled routing across the NSFNET as a whole. In fact, EGP was designed pragmatically with the hierarchical topology of the NSFNET in mind, rather than as a general purpose routing protocol for arbitrary topologies. As the original EGP specification notes, it was “intended for a set of autonomous systems which are connected in a tree, with no cycles”, and “does not enable the passing of sufficient information to prevent routing loops if cycles in the topology do exist” (Rosen, 1982, p.7). Even while EGP created new possibilities for the evolution of internet technologies, it restricted the topology of the NSFNET to a very specific form – a hierarchical structure – which has embedded within it centralised points of control at every level.

The researchers and engineers involved with the operation and development of the NSFNET were very aware of the limitations of EGP. In anticipation of the requirements of the future internet, a new routing protocol had to be devised to replace EGP, to support the more complex network topologies that might emerge in the presence of competing backbone and regional network providers. As a network administrator involved with the design and operation of the NSFNET backbone told me, “there was an urgency to convert” away from EGP, by designing and deploying a successor, the Border Gateway Protocol (BGP), “which was basically designed to avoid EGP’s shortcomings”.

The initial version of BGP was standardised in 1989 to address the problems with EGP, not the least of which was support for complex network topologies (Lougheed & Rekhter, 1989). BGP allowed loops in the topological organisation of networks, by the simple expedient of requiring each network to append its unique autonomous system number to every BGP routing announcement it originated or relayed. A network could simply ignore routing announcements in which its autonomous system number was embedded, since this indicated that the routing announcement had already transited the network once before, and was arriving as part of a loop in network topology. In theory, BGP made possible the construction of arbitrary network topologies without any need for centralised points of control.

Although BGP importantly supported future network topologies, it was designed without security requirements in mind. By accidental misconfiguration, or through failures of hardware or software, networks sometimes made claims through BGP about routes to which they could carry traffic, which in actuality they could not reach. This is not to say that routing on the NSFNET was insecure: the hierarchical topology of the NSFNET provided centralised points of control through which stable routing could be assured.

A centralised database of routing information was maintained by the operators of the NSFNET backbone, called the Policy Routing Database (PRDB). The connection of a new network to the NSFNET, or the allocation of new IP address space to an existing connected network, could only be made effective once these details were recorded in the PRDB. Updates to the PRDB required coordination – typically by email – between key administrative personnel at regional networks and the NSFNET backbone, who trusted each other to be responsible and provide accurate

routing information.

As many of my interviewees told me, "everybody trusted everybody" in the tightly knit research community developing and operating the NSFNET. Security was not an important issue within the community of the NSFNET's technical personnel, who trusted those responsible for operating centralised points of control, and relied upon interpersonal trust relationships in the everyday practice of operating the NSFNET (Mathew, 2014; Mathew & Cheshire, 2010). The apparent lack of security in BGP was as much an outcome of its operating environment – the hierarchical topology of the NSFNET – as it was of trust relationships among the NSFNET's technical personnel.

The regional networks, and the NSFNET backbone, applied filters based on PRDB information to their BGP routers, ensuring that only valid BGP routing announcements would be relayed through them. Any routing claims which did not match information in the PRDB were stopped from propagating further through the NSFNET (whether at the regional network, or the NSFNET backbone), ensuring that routing on the NSFNET remained stable and secure.

The reliable operation of BGP on the NSFNET was made possible through two kinds of centralised control. First, through the hierarchical topology of the NSFNET that provided locations at which filters could be applied to block incorrect BGP routing announcements. Second, through the central position of the NSFNET backbone that allowed its administrators to establish the practices by which the NSFNET was operated, giving them the political capacity to require the use of a centralised database for routing information.

The result was a hierarchical system of governance, mirroring the hierarchical topology of the NSFNET, with a strong centre of control in the form of the PRDB at the NSFNET backbone. The earliest routing protocol for this system, EGP, was designed for this hierarchical topology. The design of BGP, as a successor to EGP, took place in the context of this hierarchical topology, delegating responsibility for security in a potentially complex graph of interconnections to the central location of the PRDB, managed by the administrators at the NSFNET backbone.

SECTION 6: DISTRIBUTED CONTROL IN THE COMPLEX GRAPH OF THE INTERNET

The centralised controls involved in the operation of the NSFNET strongly influenced plans for the privatisation of the NSFNET, shaping the emergence of the public internet. The technical personnel involved in planning for the transition to a more complex network topology – with multiple competing backbone and regional networks – anticipated that a centralised database for routing information, which they termed the Route Arbiter (RA), would continue to be required for the reliable operation of BGP. A network administrator involved in setting up the RA told me how it was intended to serve a similar function to the PRDB on the NSFNET, as an authoritative source for routing information: "if there's any dispute for the routing, then the Routing Arbiter will decide." Accordingly, the NSF issued a solicitation for an organisation to take on the role of the RA for the internet (NSF, 1993).

The contract for building and maintaining the RA was awarded early in 1994 to The Merit Network and the University of Southern California's Information Sciences Institute. They created the Routing Assets Database (RADb), which remains in use today as a shared, publicly available store of routing information, with this information maintained voluntarily by network

administrators from autonomous systems across the internet.

Unlike the PRDB, which could be used to enforce routing policies due to its central position in the NSFNET backbone, the RADb was a third party service which had no way to require networks to use it. In fact, many networks actively chose not to put their routing information into the RADb. The reasons for this reluctance became clear in my interviews, as technical personnel involved in building networks in this period told me how they were worried about publicly exposing sensitive data about network configurations and customers. As one of my interviewees pointed out, “some of the ISPs, because of competitiveness, they don't really want other people to know who their clients are; but if you update RADb [with your routing information], it's public.”

In consequence, the routing information in the RADb was incomplete and inconsistent, making it an unreliable source of routing information. Several networks created their own parallel RADb implementations, collectively termed Internet Route Registries (IRRs). In the absence of a single reliable centralised database of routing information - to support the verification of routing claims in BGP announcements - it became substantially more difficult to maintain the secure, stable operation of network interconnection through BGP, as I describe below.

Once the NSFNET was privatised, in 1995, many private providers of internet services quickly emerged. The resultant network topology was a complex graph of connectivity, quite unlike the ordered hierarchy of the NSFNET. The earliest efforts to understand the topology of the internet focused on its aggregate properties – such as the number of interconnections for each network on the internet – observing their close conformance to power law relationships (Faloutsos, Faloutsos, & Faloutsos, 1999). Subsequent efforts attempted to examine the actual structure of internet topology, commenting on the difficulty of this task, since there was no longer a single central location (such as the NSFNET backbone) from which to observe topological form (Gao, 2001; Oliveira, Pei, Willinger, Zhang, & Zhang, 2008; Oliveira & Willinger, 2010). All knowledge of topology on the internet is necessarily partial, bounded by the locations of the networks from which topology is observed.

Even with only partial knowledge, it is readily apparent from these studies that the topology of the internet is far from decentralised. A few tens of networks occupy central positions in global internet topology, providing international connectivity spanning countries and continents. Within regions and countries, there are often dominant networks providing transit to the global internet. The resulting topological structure is composed of multiple centres of power, at different geographical scales. In general, smaller networks rely on larger networks to carry traffic to destinations across the internet, while larger networks employ the staff and tools required to manage their more highly connected topological positions, in a complex web of technical-economic-political dependencies.

It has proved difficult to reliably operate BGP over the complex graph of internet topology, and with only partial knowledge of topological structure. BGP has been subject to ongoing failures, attributable to administrative mistakes and technical issues, and also to active attacks seeking to intercept and divert internet traffic (Boothe, Hiebert, & Bush, 2006; Goldberg, 2014; Khare, Ju, & Zhang, 2012; Schlamp, Carle, & Biersack, 2013). The responses to these problems have been both social and technological.

Social trust relationships are essential to the functioning of the modern internet, just as they were for the NSFNET. Ordinary users perceive the internet as stable because of social trust relationships amongst technical personnel at networks located at key topological positions in the

internet. These trust relationships are produced within regionally organised professional communities (such as NANOG), and are leveraged in the everyday practice of network interconnection to enable coordination and collaboration in ongoing efforts to repair and mitigate routing faults propagated through BGP (Mathew, 2014; Mathew & Cheshire, 2010).

The technological effort underway at the IETF to secure BGP takes the form of the Resource Public Key Infrastructure (RPKI), an extension to BGP which provides mechanisms to allow networks to verify the routing claims made in BGP announcements. ⁴RPKI relies on centralised authorities to provide the cryptographic signatures necessary to authenticate routing claims. The deployment of RPKI is being led by the Regional Internet Registries (RIRs), which are the centralised authorities responsible for uniquely allocating IP address space and autonomous system numbers to enable the operation of BGP. Some in the internet's technical communities view RPKI with suspicion - even though they might trust the RIRs - being wary of giving up control of their routing in any way to a centralised authority. Others regard RPKI as a necessary step towards securing BGP. The future is uncertain, but it is clear that the technological choice to use (or not use) RPKI brings with it distinctive governance arrangements, with varying degrees and kinds of centralisation.

In either case, the practice of operating BGP is anchored by the centralised resource allocation functions of the RIRs, and the centralised standardisation function of the IETF; and regulated by national and international telecommunications bodies and law. Power and control in internet topology is necessarily distributed between functionally-specific centres, regionally organized technical communities, the more centrally located networks in global and regional topological structure, and national and international regulatory authorities.

CONCLUSION

As I have shown, the infrastructure of the internet has never been decentralised, nor was it designed with decentralisation as a primary goal. At every stage of the internet's history, there have been centres of control, necessary for the operation of internet infrastructure: from control through a single administrative centre in the ARPANET, to a hierarchy of control and coordination in the NSFNET, to multiple geographically distributed centres of control in the internet, to a possible future in which RPKI establishes new centres of control. Centralised governance and decentralised operation of the internet are not mutually exclusive conditions. Each is responsible for specialised functions, depending upon the other to achieve the outcome of a stable, reliable and secure internet infrastructure.

Public debates about the internet often decry the present, in which the freedoms offered by the internet are under attack by nation states and monopolistic corporations. These debates invoke a past in which the internet is said to have been freer, and more decentralised, and imagine a possible future in which freedoms might be ensured once more through the development of decentralised technologies, immune to control.

However, as I have shown, if the internet encapsulated certain freedoms in the past, it was by no means as a consequence of intrinsically decentralised technology. All systems have centres of power, whether as global administrative functions (such as centralized administrative control in the ARPANET, or the NSFNET PRDB) or as concentrations of power in topology (such as the more central networks which control routing across specific geographical scales).

The danger inherent in imagining and designing future internet technologies as though they are decentralised is that the inevitable centres of control required for governance will go unremarked upon, and in doing so will become susceptible to political capture by the very powers that these technologies seek to evade. It is critical to anticipate and design for the functions and accountability of centres of power in internet technologies. For instance, the function and accountability of centers of power may be intentionally designed, as with the function of the PRDB as the authoritative source of routing information for the NSFNET, which was held to account through social trust relationships among NSFNET network administrators. Equally, the intention behind design for function and accountability may fail, or be only partially realised, as was the case with the RADb for the public internet. In this case, the result is a functional reliance upon networks in relatively more central topological positions to maintain the stability and security of BGP; with these more central networks held to account through social trust relationships among the technical personnel managing interconnections between networks on the internet.

Restricting the terms of the debate around future internet technologies to centralised vs. decentralised limits possible visions to purely technological solutions, whether as architectures for control, or as mechanisms to escape control. The emblematic representation of the centralised vision is the panopticon, a central point of perfect observation and regulation over society (Foucault, 1975). In contrast, the decentralised vision imagines scientists and engineers with perfect knowledge of social relations, who are able and willing to devise decentralised technologies capable of supporting all of society. Paradoxically, both centralised and decentralised visions are totalising, permeating everywhere, yet located nowhere, existing independent of space and time.

These visions are representative of what Haraway (1988) termed "the view from nowhere", a perspective which presents scientific knowledge as absolute, valid everywhere and everywhen. Alternative approaches to understanding the constitution of knowledge call for attention to the limits to perspective, and the situated contexts, through which knowledge is constructed in practice (Haraway, 1988; Lave & Wenger, 1991). The internet is neither centralised nor decentralised: throughout its history, it has functioned through situated knowledges constructed from partial perspectives, dependent upon topological and administrative locations. Perhaps a new language is required to make sense of the internet as it was, is, and can be.

I propose that we think of the internet as distributed, rather than in the dichotomy between centralised and decentralised. A perspective on the internet as a distributed system acknowledges that concentrations of power are inevitable, and sometimes necessary. To enable potential freedoms on the internet imagined as distributed system, power must be dealt with on its own grounds, with solutions that are as much political as they are technological. Equally, freedom must be framed in terms of the interpersonal relationships and obligations required for the stable operation of interdependent technological systems. Technologies provide architectures for the world; but we inhabit them, shape them, and are shaped by them.

Centers of power – whether administrative or topological – must be held to account, through regulation and representation in their function. Equally, technology must be held to account, for the contingent processes through which practices, relations and institutional structures emerge to stabilize and order technological form. The question to be asked is how particular articulations of technological form, topology, and administrative structure and practice serve the public good. The answers to this question are necessarily partial and pragmatic, based in particular contexts, drawing from situated knowledges. But it is through the relations between

these partial answers that the governance of a distributed system may be comprehended.

REFERENCES

- Abbate, J. (1999). *Inventing the Internet*. MIT Press.
- Allen, J. (2011). Topological Twists: Power's Shifting Geographies. *Dialogues in Human Geography*, 1(3), 283–298. doi:10.1177/2043820611421546
- Barlow, J. P. (1996). *A Declaration of the Independence of Cyberspace*. Retrieved from <https://www.eff.org/cyberspace-independence>
- Beniger, J. R. (1986). *The Control Revolution: Technological and Economic Origins of the Information Society*. Harvard University Press.
- Benkler, Y. (2016). Degrees of Freedom, Dimensions of Power. *Daedalus*, 145(1), 18–32. doi:10.1162/DAEDa00362
- Boothe, P., Hiebert, J., & Bush, R. (2006). How Prevalent is Prefix Hijacking on the Internet? In *Proceedings of NANOG36*. Retrieved from <http://www.nanog.org/meetings/abstract?id=411>
- Bowker, G. C., & Star, S. L. (2000). *Sorting Things Out: Classification and Its Consequences*. MIT Press.
- Braman, S. (2011). Privacy by design: Networked computing, 1969-1979. *New Media & Society*, 14(5), 798–814. doi:10.1177/1461444811426741
- Coutard, O. (1999). *The Governance of Large Technical Systems*. London; New York: Routledge.
- DeNardis, L. (2009). *Protocol Politics: The Globalization of internet Governance*. MIT Press.
- Edwards, P. N. (2003). Infrastructure and Modernity: Force, Time, and Social Organization in the History of Sociotechnical Systems. In T. J. Misa, P. Brey, & A. Feenberg (Eds.), *Modernity and Technology* (pp. 185–226).
- Faloutsos, M., Faloutsos, P., & Faloutsos, C. (1999). On power-law relationships of the Internet topology. *ACM SIGCOMM Computer Communication Review*, 29(4), 251–262. doi:10.1145/316194.316229
- Foucault, M. (1975). *Discipline & Punish: The Birth of the Prison*. Vintage.
- Galloway, A. (2006). *Protocol: How Control Exists after Decentralisation*. MIT Press.
- Gao, L. (2001). On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Transactions on Networking*, 9(6), 733–745. doi:10.1109/90.974527
- Goldberg, S. (2014). Why Is It Taking So Long To Secure Internet Routing? *Communications of the ACM*, 57(10), 45–52. doi:10.1080/1357628021000012705
- Granick, J. (2015). *The Death of the Internet Dream*. Backchannel. Retrieved from <https://medium.com/backchannel/the-end-of-the-internet-dream-ba060b17da61>
- Haraway, D. (1988). Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. *Feminist Studies*, 14(3), 575–599. Retrieved from <http://www.jstor.org/stable/3178066>

Harris, S. R., & Gerich, E. (1996). Retiring the NSFNET Backbone Service: Chronicling the End of an Era. *Connexions*, 10(4). Retrieved from

http://www.merit.edu/networkresearch/projecthistory/nsfnet/nsfnet_article.php

Harvey, P. (2012). The Topological Quality of Infrastructural Relation: An Ethnographic Approach. *Theory, Culture & Society*, 29(4-5), 76–92. doi:10.1177/0263276412448827

Heart, F. E. (1969). *Interface Message Processors for the ARPA Computer Network*. Retrieved from <http://www.dtic.mil/dtic/tr/fulltext/u2/686811.pdf>

Hinden, R., & Sheltzer, A. (1982). *RFC 823: The DARPA Internet Gateway*. RFC Editor. Retrieved from <https://tools.ietf.org/html/rfc823>

Hofmann, J., Katzenbach, C., & Gollatz, K. (2016). Between coordination and regulation: Finding the governance in Internet governance. *New Media & Society*. doi:10.1177/1461444816639975

Hughes, T. P. (1983). *Networks of Power: Electrification in Western Society, 1880-1930*. Johns Hopkins University Press.

Hughes, T. P. (1987). The Evolution of Large Technological Systems. In W. E. Bijker, T. P. Hughes, & T. J. Pinch (Eds.), *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology* (pp. 51–82). MIT Press.

Khare, V., Ju, Q., & Zhang, B. (2012). Concurrent Prefix Hijacks: Occurrence and Impacts. In *Proceedings of the 2012 ACM Conference on Internet Measurement* (pp. 29–35). Retrieved from <http://dl.acm.org/citation.cfm?id=2398780>

Lave, J., & Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press.

Lessig, L. (2001). *The Future of Ideas: The Fate of the Commons in a Connected World*. New York: Random House. Retrieved from <http://www.the-future-of-ideas.com/>

Lougheed, K., & Rekhter, Y. (1989). RFC 1105: A Border Gateway Protocol. RFC Editor. Retrieved from <http://www.ietf.org/rfc/rfc1105.txt>

Mathew, A. J. (2014). *Where in the World is the Internet? Locating Political Power in Internet Infrastructure*. University of California, Berkeley. Retrieved from <http://www.ischool.berkeley.edu/research/publications/ashwinmathew/2014/whereworldinternetlocatingpoliticalpowerinternetinfrastructure>

Mathew, A. J., & Cheshire, C. (2010). The New Cartographers: Trust and Social Order within the Internet Infrastructure. In *Proceedings of the 38th Research Conference on Communication, Information and Internet Policy (Telecommunications and Policy Research Conference)*. George Mason University School of Law, Arlington, VA. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1988216

Mayntz, R., & Hughes, T. P. (Eds.). (1988). *The Development of Large Technical Systems*. Campus Verlag; Westview Press.

Mueller, M. L. (2010). *Networks and States: The Global Politics of Internet Governance*. MIT Press.

- Mumford, L. (1964). Authoritarian and Democratic Technics. *Technology and Culture*, 5(1), 1–8. Retrieved from <http://www.jstor.org/stable/3101118>
- 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*. Retrieved from http://w2.eff.org/Infrastructure/NRENNSFNETNPN/nsf_nren.rfp
- Oliveira, R., Pei, D., Willinger, W., Zhang, B., & Zhang, L. (2008). In search of the elusive ground truth: the Internet's AS-level connectivity structure. *SIGMETRICS Perform. Eval. Rev.*, 36(1), 217–228. doi:10.1145/1384529.1375482
- Oliveira, R., & Willinger, W. (2010). The (In)Completeness of the Observed Internet AS-level Structure. *IEEE/ACM Transactions on Networking*, 18(1), 109–122. doi:10.1109/TNET.2009.2020798
- Rosen, E. C. (1982). *RFC 827: Exterior Gateway Protocol*. RFC Editor. Retrieved from <http://tools.ietf.org/html/rfc827>
- Schlamp, J., Carle, G., & Biersack, E. W. (2013). A Forensic Case Study on AS Hijacking: The Attacker's Perspective. *ACM SIGCOMM Computer Communication Review*, 43(2), 6–11.
- Star, S. L. (1999). The Ethnography of Infrastructure. *American Behavioral Scientist*, 43(3), 377–391. doi:10.1177/00027649921955326
- Star, S. L., & Ruhleder, K. (1996). Steps Toward an Ecology of Infrastructure: Design and Access for Large Information Spaces. *Information Systems Research*, 7(1), 111–134.
- Strazisar, V. (1979). *IEN 109: How to Build a Gateway*. Retrieved from <http://www.postel.org/ien/pdf/ien109.pdf>
- Strazisar, V., & Perlman, R. (1978). *IEN 30: Gateway Routing: An Implementation Specification*. Retrieved from <http://www.postel.org/ien/pdf/ien030.pdf>
- van Eeten, M. J., & Mueller, M. (2013). Where is the governance in Internet governance? *New Media & Society*, 15(5), 720–736. doi:10.1177/1461444812462850
- Winner, L. (1980). Do Artifacts Have Politics? *Daedalus*, 109(1), 121–136. Retrieved from <http://www.jstor.org/stable/20024652>

FOOTNOTES

1. In general, I use the term “network” to refer to the individual computer networks which make up the internet, as opposed to a network connecting people or things
2. IETF standards documents are available at <https://ietf.org/rfc.html> IENs are archived on the website of the Postel Center, <http://www.postel.org/>
3. For a current count of autonomous systems visible in the inter-domain routing system, see the entry “Number of ASes in routing system” on this report: <http://www.cidr-report.org/as2.0/>
4. For more information on RPKI, see the work of the IETF Secure Inter-Domain Routing Working Group: <https://datatracker.ietf.org/wg/sidr/documents/>.