

'This is not how we imagined it': Technological affordances, economic drivers and the Internet architecture imaginary

Niels ten Oever, University of Amsterdam, the Netherlands

Abstract

The Internet architecture is widely perceived as engine for innovation by providing the equal opportunity to deploy new protocols and applications. This view reflects an imaginary that guides the co-production of policy and technology that can be traced back to the early days of the Internet, which is still prominent among the engineers in one of the main governance bodies of the Internet: the Internet Engineering Taskforce. After the privatization of the Internet architecture in the 1990s, the interplay between the architectural principles of end-to-end, permissionless innovation, and openness subverted equality among Internet users and hampered their ability to redesign the Internet. I draw on media studies, science and technology studies and international political economy, and use a combination of qualitative and quantitative methods to show how the Internet architecture's affordance structure got reconfigured, and how this facilitated the prioritization of corporate interests over the interests of end-users.

Keywords:

Internet governance, standard setting, sociotechnical imaginaries, end-to-end, permissionless innovation, architecture, infrastructure, media studies, STS, governance studies

Corresponding author: Niels ten Oever, Singel 40 F, 1015 AB Amsterdam, The Netherlands. Email: mail@nielstenoever.net

Introduction

When in the early nineties the Internet was released from the labs and found its way to millions of users, it was widely perceived as an engine for innovation (Van Schewick, 2012), an information highway (Flichy, 2007) and a tool for democratization (Castells, 2009). These expectations and aspirations accompanied the development of the Internet architecture and were operationalized through three main architectural principles: end-to-end, permissionless innovation, and openness. In this paper, I show how the interplay between these principles, after the privatization of the Internet in the early 1990s, undermined the equality of users and the ability of individuals, researchers, and small companies to redesign the Internet.

The Internet architecture is co-produced by corporations, state actors, researchers and advocates in a self-regulatory industry standards body called the Internet Engineering Taskforce (IETF).¹ Self-regulation has been the general paradigm for the governance of the Internet because it is assumed to be most suited to the transnational and quickly evolving nature of the Internet (Price and Verhulst 2000). The evolution of the architecture of the Internet is taking place through the development of open and voluntary standards that facilitate interoperability between the products of network operators, equipment vendors, content and service providers, and software developers. Because of the nature of the standard setting process, it has been described as a 'wild mix of politics and economics' (Shapiro and Varian, 1998) and 'politics by other means' (Abbate, 1999). While the standards and protocols that are developed in the IETF are largely hidden from the larger public, they shape our behavior

¹ The IETF is by no means the only standards body involving the Internet. For instance the World Wide Web Consortium (W3C) sets standards for the web. Other examples are the Third Generation Partnership Project (3GPP) and the Institute for Electrical and Electronics Engineers (IEEE) who set standards for wireless and wired communication.

(Chadwick, 2006), determine vectors of control over user data flows (Galloway, 2006), how users access information (DeNardis, 2014), how users can exercise their rights online (Lessig 2006).

To understand the standard setting process I use the terms ‘sociotechnical imaginary’, ‘co-production’, and ‘technological affordance’. A sociotechnical imaginary is the combination of visions, symbols, and futures that exist in groups and society. It influences behavior, individual and collective identity as well as the development of narratives, policy, and institutions (Jasanoff and Kim, 2015). A sociotechnical imaginary guides the process in which people co-create knowledge, technology, and order, a process that Jasanoff (2004) calls co-production. Technology, which is an inherent part of this co-production process, inhibits and stimulates human behavior. Hutchby (2001: 441) describes this ‘constraining, as well as enabling, materiality of artifacts’ as technological affordances. I leverage these terms to show how the Internet architecture’s sociotechnical imaginary and its technological affordances got reconfigured and subverted during three decades of co-production following the commercialization and privatization of the Internet. This compounded theoretical lens allows me to simultaneously take into account the shaping of institutional configurations, technological orderings, economic drivers, and the collaboration among disparate groups and competitors facilitated by a joint vision. This approach enables me to analyze the Internet architecture as a site of contestation (ten Oever, 2019), as an assemblage of power (DeNardis, 2014), and ‘as a normative "system of systems"’, and to unpack ‘the micro practices of governance as mechanisms of distributed, semi-formal or reflexive coordination, private ordering, and use of internet resources’ (Epstein et al., 2016), without defaulting to a reductionist approach.

First, I will provide an overview of the relevant literature, subsequently I give an overview of the methods used in this research, and provide an analysis in which I will establish the Internet architecture imaginary and subsequently show how it got subverted. Finally, I will offer some thought about what this means for self-regulatory Internet governance models and avenues for further research.

An imaginary space between a technological dream and an economic reality

Instead of looking at the content of data streams, which is like the ‘juicy piece of meat carried by the burglar to distract the watchdog of the mind’ (McLuhan 2013, 19), I argue we should rather look at the *preconditions, shapes, and characteristics* of data streams: the Internet architecture. Before elaborating this position, I will first describe the process through which the Internet architecture is co-produced and introduce the concept of the sociotechnical imaginaries as a lens to understand this process and give an overview of recent academic debates pertaining the Internet architecture. Contemporary debates in media studies, science and technology studies and governance studies that discuss the Internet architecture focus on: (1) the values, or lack thereof, that are enshrined in the internetworking protocols (Braman, 2012b; Flanagin et al., 2009), (2) how the Internet infrastructure is used to exercise control (Eeten and Mueller 2013; Musiani et al. 2015), and (3) consolidation and market concentration in the Internet architecture (Easterling, 2014; Mansell, 2013; McKelvey, 2018).

Technical standards, of which networking protocols are a subset, are rules, procedures and formats that facilitate communication between two or more parties. The Internet architecture consists of ‘standards which make up the technical back-bone of an information infrastructure’ (Hanseth and Monteiro, 1997: 183) that, through its affordance structure, dynamically shapes our information societies. The Internet architecture is co-produced in governance bodies and Standards Developing Organizations such as the

IETF. Whereas theoretically participation in the IETF is open for everyone, it is dominated by employees of transnational corporations. The most common affiliations of the authors of IETF output documents, the so-called RFC-series, are: Cisco, Huawei, Ericsson, Google, Juniper, IBM, Nokia, Microsoft, AT&T, and BBN.² The RFC-series (an abbreviation for Request for Comments) should not only be understood as a series of technical documents but also as policy documents (Braman, 2013) which describe the values, processes and procedures for co-production and therefore are relevant for understanding the sociotechnical imaginary of the Internet architecture. Most RFCs, however, are written by authors with an affiliation in the private sector, there are also many RFCs that have been authored by researchers, and even some by members of civil society organizations.

The Internet architecture's sociotechnical imaginary revolves around doing things that are 'for the good of the Internet' (Mathew, 2014: 160), sustaining a 'generative Internet' (Zittrain, 2008: 6), and is underpinned by three specific engineering principles: end-to-end, permissionless innovation, and openness (Internet Society 2012). The Internet architecture's imaginary is rooted in the idea that the network is a general purpose 'common carrier network' (Davies et al., 1967: 3) where 'all hosts are equal' (Mogul et al., 1984: 1), meaning that they can function as general purpose end nodes (cf. Padlipsky 1982; Braden 1989; Carpenter 1996), and that everyone has the freedom to shape their own traffic by deploying new protocols between end nodes, and thus redesigning the Internet. The ability to freely deploy protocols fits very well with the 'ideology of open standards' (Russell, 2014: 21) and the voluntary nature of the Internet standards developed in the IETF. In this paper I interrogate this sociotechnical imaginary because '[m]yths are important for what they reveal (including a genuine desire for community and democracy) and for what they conceal (including the growing concentration of communication power in a handful of transnational media businesses)' (Mosco, 2005: 19).

DeNardis (2009) explores the complex process of the co-production of the Internet architecture by describing how the mere transition from one protocol to another caused a significant amount of contestation because of its geopolitical interests and impacts. This process of contestation is described by Clark et al (2005), who argue that there are different adverse interests at work in defining the Internet architecture, that this 'tussle' should be accommodated because 'it is crucial to the evolution of the network's technical architecture' (Clark et al. 2005, 65), and that rigid designs which do not accommodate this tussle will not survive the passage of time and will be broken. Braman (2011) convincingly shows that social policy issues such as rights and freedoms have always been part of the Internet standards deliberations. Braman (ibid), however, did not address how or whether these discussions actually resulted in changes in the technical materiality of the network. Davidson, Morris, and Courtney foresaw Braman's findings but argue that while 'many technologists within the leading standards bodies are public-minded, few have explicit expertise in policy-making or at interpreting the public interest. Standards organizations have always (appropriately) emphasized technical goals over societal ones' (Davidson et al., 2002: 4). Since the call of Davidson et al to assess the impact of protocols in the IETF there have been several efforts to better understand the relationship between values and networks (Orwat and Bless, 2016), develop guidelines to integrate human rights considerations in protocol design (ten Oever and Cath, 2017), and calls for the IETF to 'enable the actualisation of human rights through the protocols and standards it designs by implementing a responsibility-by-design approach to engineering' (Cath and Floridi, 2017: 449).

² <https://www.arkko.com/tools/allstats/companydistr.html> accessed on November 17, 2018

The IETF has not operationalized any structural assessment of the impacts of their standards and protocols. The lack of integration of impact assessments in the standards process and the intentional undermining of technical standards has led to a discussion about the legitimacy and adequacy of the self-governing technical standards bodies to deliver a trustworthy Internet architecture (Rogers and Eden, 2017). Internet shutdowns during political events, such as elections, foregrounded how the Internet architecture is used as a domain of control and showed how infrastructure is used by governments to realize their objectives. Research has established a trend of the enactment of governance objectives through and by private parties rather than governments (Arpagian, 2016; DeNardis and Musiani, 2014). Levinson and Cogburn (2016: 219) remark that this process is tightly connected with the privatization of the governance of the Internet architecture. While the privatization of the Internet architecture was supposed to lead to competition and innovation (Cowhey et al., 2009; Van Schewick, 2012), this paper argues that actually led to the subversion of equality between hosts and the freedom to deploy new protocols.

As described above the Internet architecture has been an object of research in different fields, but analyses that take into account the combination of the guiding sociotechnical Internet architecture imaginary, the materiality of the technology, and economic drivers, are still quite rare, but are gaining traction (eg. DeNardis 2014; Mathew 2014; McKelvey 2018). Analyses generally take into account institutional and technical aspects, such as Dourish (2018) has done in his analysis of IPv6; or rather, it takes economic and institutional aspects into account, such as Van Schewick (2012), Russell (2014) and Smyrniotis (2018) have done. Arguably, both approaches undervalue the dynamic interplay between technological materiality, institutional configurations, and economic drivers. I argue that it is exactly this interplay that creates new orderings and affordances. My theoretical contribution is to reveal how economic drivers prompted an interplay among architectural principles which led to a reconfiguration and subversion of technological affordances and the Internet architecture's sociotechnical imaginary. This seeks to overcome a economic, legal, or technological reductionist approach in the analysis of the Internet architecture.

Methods

My research into the IETF started with a long-standing fascination for RFCs: their language, particular formatting, and authoritative standing for everyone interested in computer networking. The institutions, people, and processes behind the production of the RFCs, their infrastructure so to say, only became apparent when I started participating in the IETF and its surrounding environs. This participation is the basis for my ethnographic memoir, which developed into participant observation when I formalized my research plans. The research period spanned between March 2014 and July 2018, during which I participated in 11 tri-annual IETF meetings and actively participated on mailinglists. I participated in several Working Groups, and served in several leadership positions. This experience provided me with a first-hand account of the practices in the most prominent Internet standards body, as well as access that an external observer might not otherwise have. In qualitative research, the researcher is an inherent part of the creation of meaning (Denzin et al., 2006), part of an critical ethnographic practice is therefore 'an ongoing commitment to re-thinking and re-doing one's work as ethnographer and activist' (Lave, 2011: 2). Part of this process was to address my particular *situatedness* in the fieldsite (Haraway 1988), namely as an activist-engineer-researcher. In order to gather and seek to understand different points of views, I employed a mixed method approach, to triangulate and validate my findings, and in that process to create an opportunity for reflection on research context, the relationships with the

community I researched and was situated in, and the power dynamics in the process of knowledge production (Haraway 1991).

In order to analyze the evolution and emergence of explicit values in the large body of data on transnational governance of the Internet infrastructure, I engaged in the quantitative analysis of IETF mailinglists and the IETF's technical documents published in the so-called RFC-series. In my analysis I focused on prevalence and development over time of language related to society, ethics and rights, as well as trends in the professional affiliation of document authors, guided by intuitions that arose from participant observation. I obtained these documents from the IETF website, after which I undertook a quantitative document analysis of the RFC-series using the JavaScript-based tool *rfc-analysis*³, and engaged in quantitative mailing list analysis using the Python-based tool *BigBang*⁴, to gain a deeper insight into cases, trends and interactions. The outcome of these quantitative analyses informed the creation of the questionnaire I developed for the interviews, and helped me focus on architectural principles and specific protocols. I engaged in 25 semi-structured interviews with IETF leadership and RFC authors through a purposive sample of seasoned and visible members of the Internet protocol community. The audio was transcribed and analyzed using qualitative methods informed by thematic analysis, with which themes were identified and coded across interviews. Through the identification of themes, concepts, practices, and activities, I analyzed the interview data to understand the ways in which the Internet architecture fundamentally changed from the early 1990s up to now, and how that affected the equality of users and their ability to design and deploy new protocols. The last step of triangulation and validation was to see how my findings on the Internet architecture compared to the description of the Internet architecture in the RFC-series. Therefore I engaged in the qualitative analysis of a specific subset of RFCs. I made a purposive sample of 20 RFCs that mention 'architecture' and made a chain sample of 20 RFCs that are referred to in the aforementioned 20 RFCs to add background understanding for the architectural issues that are being referred to and that they are responding to. Finally, I analyzed 5 RFCs that specifically got mentioned during the interviews.

Analysis - On the idea of smart endpoints and the dumb pipes

I will first describe the Internet architecture imaginary, then I will describe the challenges in the form of the rise of the middlebox and the accompanying reconfiguration of the affordances of the architecture, and the subsequent iterative responses to overcome the obstacles presented by this new ordering.

The Internet architecture imaginary

The sociotechnical Internet architecture imaginary emerged during the early phases of the development of the Internet while it was still a research network. I focus on the stabilization of this imaginary that started with the privatization of the Internet in the early nineties when the US government ceded direct control over the Internet. This gave way to an increased amount of self-regulation through private governance bodies such as the IETF. When asking engineers in the IETF about the central architectural values or principles of the Internet protocol community, their answers have a significant amount of overlap. I will describe the imaginary as a category, or ideal type, based on the research data, which in reality can appear less monolithic and will have fuzzy edges. Nonetheless in the interviews, documents, and observations, the imaginary turned out to be remarkably consistent. The end-to-end principle (Saltzer, Reed, and Clark 1984), permissionless innovation, and openness (Russell 2014) get mentioned

³ <https://github.com/npdoty/rfc-analysis> access on October 11, 2018

⁴ <http://dataactive.github.io/bigbang/> accessed on August 25, 2018

time and again in interviews as well as in technical and policy documents (Internet Society, 2012). These three architectural principles shaped a sociotechnical imaginary which is rooted in the equality of ‘internet host computers’ (Deering, 1989; Mogul et al., 1984), the ability to design and deploy new protocols between these computers, and to increase and grow the Internet with more computers and more users (Braman, 2012a). The architectural principles have both sociotechnical and sociopolitical conceptions which play important roles in the co-production of the Internet. I will discuss the three architectural principles in depth, because they played a central role in the demise of the Internet architecture’s imaginary.

The first architectural principle, the end-to-end principle, appeared as a central pillar of the values and principles of the architecture in nearly every interview, RFC3724 even calls it: ‘the core architectural guideline of the Internet’ (Internet Architecture Board, 2004). The principle describes where to put functionality in the network, namely at the edges (Internet Architecture Board, 1996), and let the network be ‘dumb pipes’⁵ that solely transports data. The end-to-end principle allowed for a ‘tremendous amount of agency in individuals and anyone who could put a server anywhere. Anybody could make arrangements to have a prototype protocol pair that you could talk to with each other from anywhere to anywhere.’⁶ This principle was infra-structurally a revolution because it contrasted so strongly with the communication networks that preceded the Internet. Endpoints that are controlled by their owners can be altered quickly, and thus allow for freedom and flexibility. Changes in the infrastructure are far more cumbersome, or in the words of a former telecommunications engineer turned Internet engineer: ‘we have the end-to-end principle because so you can do things really quickly on the infrastructure, but [...] if you have to change the infrastructure, that takes a long time’⁷. This captures the importance of the end-to-end principle for innovation, but it has further implications as a sociopolitical conceptualization, one engineer mentioned that:

There are other folks who take that principle to be more than an engineering principle, but rather an ethical or values driven principle which says that the role of the network is to enable parties to communicate with each other, and not to enable the network itself as a form of control, centralized control. I take both views.⁸

The end-to-end principle provides users with the freedom to shape and create their own networking experience. This had a tremendously empowering effect on engineers: there was ‘a desire to go your own way, um, and a kind of idea that we can invent our own rules and we don’t need too many rules, but the ones that we want, we can invent’⁹.

The second principle, the principle of permissionless innovation, describes that there should be no barriers to the deployment of new protocols. In other words: ‘you don’t need to negotiate with any entity in the middle of the network to get your new thing deployed. [...] [Y]ou don’t need to negotiate with any entity in the middle of the network in order to transport your packets.’¹⁰ Negotiating is meant here in both technical as well as sociopolitical terms; permissionless innovation depends on the fact

⁵ N1418 (For reasons of anonymity the names of the interviewees are not listed here. Each interviewee was coded and is distinguished from other interviewees by a number)

⁶ N0218

⁷ N0918

⁸ N0618

⁹ N1518

¹⁰ N0118

that there is no authority that can sanction what protocols can or should be used. '[A] typical example is the Web. Tim Berners-Lee did not ask permission from anyone, he invented something, went back, built it, and then it was downloaded and no one [...] had anything to say about it.'¹¹ This sociotechnical conceptualization has clear sociopolitical implications: there should be no limitations on the ability to deploy new technologies on the network. With this principle participants in the IETF also limit their own authority and responsibility, summarized in an often repeated phrase among long time participants: 'we're not the protocol police'¹².

The third principle, the principle of openness, is described as the property 'that you can reach from any point of the Internet to any other point of the Internet without your packets been hampered or they'd been stopped or so on'¹³, it furthermore means that new computers can be added to the network. The sociotechnical conception of openness is directly coupled with a conception of connectivity, access, as well as their explicit sociopolitical consequences. Except for the socio-technical conception of openness of the network, openness is also often associated with the socio-political consensus approach to standards development, which fits into an 'ideology of open standards' (Russell 2014, 21) that 'linked the open standards-making process with the ideals of participatory democracy, open markets, individual autonomy, and social progress' (Rogers and Eden, 2017: 804). Similar to the end-to-end principle and permissionless innovation, openness is not only associated with a technical ability, in this case to add nodes to the network, but also with open communications, open standards, as well as with open governance (Internet Society, 2013). The ideal of participatory democracy is also reflected in the IETF's unofficial motto: '*We reject: kings, presidents, and voting. We believe in: rough consensus and running code*' (Clark, 1992: 543), a credo which was minted during the Internet-Open Systems Interconnection (OSI) standards war (DeNardis, 2009; Russell, 2006), when the governance model of the IETF was heavily tested and further refined. This process strengthened the organizational practice that individuals opinions should seriously be considered and discussed, and cannot simply be overruled by an authority or a majority (Resnick, 2014). This shows the strong interrelation between the technology, the institutional organization of the IETF, and the community that co-produces the architecture. The make up of the community participating in the IETF, however, has changed over the years. In the early days of its work, the Internet architecture was produced largely by network researchers that were working at universities and as government contractors. Since the privatization of the Internet in the early nineties, there has been an exodus of researchers (Ding et al., 2013), whose ranks have been filled by contributors from the private sector who now dominate the IETF. This can, for instance, be observed in increasing preeminence of private sector affiliations among authors of RFCs.¹⁴

While the IETF community makes explicit statements about values and principles, its website says: 'We try to avoid policy and business questions, as much as possible'¹⁵. This is quite a remarkable statement for self-regulatory body of a \$44 billion networking-infrastructure market¹⁶. Interestingly, the

¹¹ N1618

¹² N0218

¹³ N0918

¹⁴ <https://www.arkko.com/tools/allstats/companydistr.html> accessed on November 19, 2018

¹⁵ <https://www.ietf.org/about/participate/get-started/> accessed on November 19, 2018

¹⁶ <https://www.srgresearch.com/articles/switch-router-revenues-set-new-record-cisco-market-share-still-over-50> accessed on March 5, 2019

architectural principles that have strong sociotechnical and sociopolitical conceptions at the same time anchor the architectural imaginary and obfuscate the socioeconomic reality.

Cracks in the Imaginary I: Firewalls, NATs, and network management

The first threat to the end-to-end principle, the openness of the Internet, and permission innovation took the form of *middleboxes*. Middlebox is a shorthand for: ‘intermediary device[s] performing functions other than the normal, standard functions of an IP router on the datagram path between a source host and destination host’ (Carpenter and Brim, 2002). Middleboxes can have many different functions, such as firewalls, Network Address Translation (NAT) routers, IP-tunnels (such as Virtual Private Networks), and network management devices. The introduction of middleboxes formed a paradigmatic shift in the functioning of the network (McKelvey, 2018). Whereas the network was previously supposed to function as a ‘dumb pipes’¹⁷, as outlined by the end-to-end principle, functionality was added to the network. This happened because of three issues that resulted from the rapid growth of the network: (1) an increased need for security, (2) the depletion of Internet Protocol (IP) addresses, and (3) increasing business interests (Internet Architecture Board, 2004). I will shortly describe the reactions to these issues below.

To be able to connect to the network every device needs a unique number: an IP address. It was never envisaged that so many devices would be connected to the network, so when more devices were connected, IP addresses were running out and a new addressing scheme needed to be developed. This was especially pressing since adding new nodes to the network is an inherent part of the principles of openness, one of the Internet’s architectural principles. However, there was no direct replacement addressing scheme ready, and there were projections that IP addresses would run out by 1994¹⁸. This led to the introduction of the ‘temporary solution’ of Network Address Translation (NAT), which allowed a network of computers to share one public IP address (Francis and Egevang, 1994) from the pool of IP addresses. While this was an efficient short-term solution, this directly went against end-to-end principle according to which ‘packets [should] flow unaltered through the network’ (Carpenter 2000). NATs interrupted the packet flow because the IP address of the end-device is not known to the network or the recipient and needed to be added by the middlebox, thus adding functionality to the network.

When the network grew beyond a group of researchers, there was the need to introduce firewalls ‘which screens network traffic in some way, blocking traffic it believes to be inappropriate, dangerous, or both’ (Freed, 2000: 2). Firewalls were installed on end-devices, home routers, but also inside larger networks and thus not only found at the edges of the Internet. A regularly implemented functionality of firewalls is directionality. This means that a network, and the computers connected to it, are ‘protected’ from receiving connections from an outside computer that it did not request. This is a sound security measure on the one hand, but on the other hand, it creates a difference between servers and clients. If your computer is located behind a directional firewall (or NAT), the computer cannot function as a server because other clients cannot reach you, traffic can only be initiated from one end of the connection. While many smartphones currently have more processor capacity and storage space than early web servers, they cannot function as a server because the network is imposing a one-directional ordering. This is how NATs and firewalls create the difference between producers and consumers.

¹⁷ N1018

¹⁸ N0718

Network operators, with the help of equipment vendors, inscribed boundaries into the Internet architecture and attempted ‘thereby to configure the user such that s/he can only meaningfully encounter the technology on the company’s terms’ (Hutchby, 2001: 451). This represents the first step in the creation of inequality between Internet hosts, and thus creation of a class of mere users.

Network management is used by network operators to optimize network performance. There are different ways for approaching network management, a contested approach is the differentiation and prioritization of specific services over others, or even the blocking of some services or providers for economic reasons. This discussion is more commonly known as the net neutrality debate (Crowcroft, 2007). One probably would not notice if you would receive an email a few milliseconds later, but if there is a delay in the delivery of a videostream, this might cause irritating hiccups. It might seem efficient to prioritize video content over mail traffic, there is, however, a fine line between network management and discrimination between kinds of traffic. If one prioritizes a specific kind of traffic or traffic from a specific provider, this could pose a barrier for alternative streams and providers to grow and develop, since competitors would have a distinct advantage. Violations of network neutrality also interfere with the end-to-end principle and the idea that the network should just transport packets.

The introduction of middleboxes in the network solved some immediate problems such as security issues, delayed other problems, such as the shortage of IP addresses, and create some economic incentives, in the case of the prioritization of services. The response to the issues of security and the lack of IP addresses could also be understood as response to the architectural principle of openness because if these issues would not be addressed it would hamper the connectivity of existing and new nodes. The changes in network management could be interpreted as enactments of permissionless innovation, but all responses inherently violated the end-to-end principle.

Cracks in the Imaginary II: The Advent of Ossification and the failure of SCTP

While middleboxes improved performance of specific kinds of traffic, they also negatively impacted the ability to alter protocols through a process called ossification (Thaler, 2011). Ossification is the decreasing flexibility of the network which results in the inability to deploy a new protocol or protocol extensions due to the unchangeable nature of infrastructure components that have come to rely on a particular feature of the current protocols (Clark 2018). NATs and firewalls ossify around specific protocol characteristics. If these middleboxes receive traffic with other, and thus unknown, characteristics they will reject the traffic. While middleboxes seek to optimize the network, they actually hamper the ability to deploy new protocols. Or in the words of a senior network operator:

So at the moment there’s a whole industry of middleboxes that basically break [...] end-to-end connections. [T]hey end up ossifying the internet itself [...] because these are boxes that are trying to operate transparently and sort of invisibly you don't know that they exist or where they exist. You can't point to them even. They don't have an address. You can't do anything. They are bumps in the wire¹⁹

Actually, ossification by middleboxes sometimes turns out to be a lot more than a proverbial bump in the wire by actually obstructing the deployment of new protocols as the following example shows.

¹⁹ N2218

The Stream Control Transmission Protocol (SCTP) was developed as an evolution to transport more data in a faster way than was possible up to then. Initially it was standardized for telephone networks in 2000 (Taylor et al. 2000) and was adapted to be a general purpose Internet protocol in 2004 and after that has received updates for over a decade. Nonetheless, SCTP never significantly worked on the Internet. SCTP worked perfectly in the lab and lived up to all of its design expectations, but it would not work *in the wild*, on the actual Internet, because middleboxes added inflexibility to the network, in other words: *ossification*. In the words of a former SCTP developer:

[Y]ou can run [SCTP] on your own network when you control all of the middleboxes, but if you try to run it across the public internet there's some non-trivial points that the traffic won't get through because there will be some boxes like: "SCTP, what's that?". NAT middleboxes are a classic example there. [...] [Y]ou can't really run SCTP across the public network. We tried that and there's too many things in the way.²⁰

Middlebox induced ossification changed the Internet from an environment where equal hosts could deploy their own protocols, to a network where to design for the future, protocols need to look like the past. Foundational architectural principles of the Internet imaginary cannibalized themselves: in order to safeguard openness, permissionless innovation in the network was leveraged. This undercut the end-to-end principle, which in turn undermined permissionless innovation.

The introduction of middleboxes reconfigured the affordances of the network, with pivoted the locus of control from the end-points to the network operators (Minar and Hedlund, 2001). The latter were enabled in this endeavor by networking equipment vendors. There were clear incentives for both the network operators and the equipment vendors: the network operators wanted more control over their networks, and offer better performance to their customers. Equipment vendors wanted to sell the network operators equipment. The way they did this was adding more intelligence to the network, which was a relatively low investment for the operators which yielded results on the short term, and benefited the network equipment vendors. An Internet pioneer who was on the forefront of connecting new countries and continents formulated it this way:

There seems to have been the development that there is now more, some would say, 'intelligence' in the network now. Well, this is a bunch of shit from a bunch of basket cases like Cisco with a willing set of co-conspirators called network operators because in the telephone world they were the center of the universe. [...] The network folks looked at this [the Internet] and said, no, no, and they found a willing co-conspirator in Cisco and instead of having 15 line router that just switched packets, now they have something with 50,000,000 lines of code.²¹

Freedom, agency, and control was taken from the endpoints by network operators, with devices that were provided by equipment vendors. As I have shown this had both technical and economic reasons, which jointly surmounted to a reordering of the affordances of the network. This reordering largely benefited network operators and equipment vendors, not so much the people that were operating services on the endpoints, because they were hampered in the deployment of new protocols, such as SCTP. Thus a response from the latter group was to be expected.

The return of the strong endpoints: The Rise of QUIC

²⁰ N2318

²¹ N0818

The limitations introduced by middleboxes accrued quite some resentment in the Internet protocol community because it constricted the freedom to deploy new protocols in the network. The fact that middleboxes do not announce themselves, and thus make the troubleshooting issues harder, added to the frustration. For quite some time, protocol developers could not find a solution: SCTP developers had worked on it for almost a decade and did not solve it. For content providers it became increasingly important to have a protocol that would deliver content in a faster manner over the networks, because of the increasing demand in streaming video and media rich websites. This finally became possible with the development of the Quick UDP Internet Connections (QUIC) protocol: a connection based stream protocol which supports multiple streams. QUIC functioned in a way similar to SCTP, with some extra features. A quintessential difference between SCTP and QUIC however was that the latter was developed by Google. Google already had a fast global Content Distribution Network and developed the most-used browser in the world: Google Chrome. Thus Google held two important pieces of the puzzle but needed a protocol to connect the two pieces: ‘Google is very invested in this [QUIC] because they make a lot of money off of making sure that no one gets in the path between them and the user, and they centralize all that power.’²². QUIC would allow Google to serve their content faster, and ensure that user data would not be shared with other parties, such as network operators. Both have significant economic implications for a company that makes most of its money via targeted advertising. But except for motive, Google also had the network control and capacity to develop this. In the words of a long time protocol developer: ‘the reason that QUIC [..] can do what it can do is because the two endpoints are controlled by the same people, so they [Google] can, they can do like dark releases and AB-testing and all that that we can't do’²³.

Google started developing QUIC in 2012 and in January 2018 between 2.1% and 9.1% of all Internet traffic was using QUIC, which is dominated by Google that uses it for 42.1% of its traffic (Rüth et al., 2018). Google was able to gain much better results than SCTP because it could do testing between its network and its browsers, and because it had significant resources to invest. Network operators would also think twice about blocking Google’s faster services because it would negatively impact many of their users:

Google’s big enough that it’s very hard to stop in the sense that when you think about blocking Google you’re blocking access to search and peoples’ email and all of the different services that they provide, a huge number of different services.²⁴

This was another large non-technical but rather economic difference between QUIC and SCTP – QUIC already had a large market share through its user-base: Google’s users. Google did not keep QUIC for itself as a proprietary protocol, Google brought QUIC to the IETF for standardization, which would increase the chances of broader adoption of the protocol, and therefore further ensure that new and updated middleboxes would not block QUIC traffic.

The implementation of QUIC will lead to a reordering of the network. QUIC was built to penetrate middleboxes and provide as little control as possible for network operators to shape, filter, or access data streams. QUIC was built to reconfigure the affordances of the architecture: it will reinstate the

²² N2118

²³ N0218

²⁴ N2118

end-to-end principle and re-enable permissionless innovation but only as long as the QUIC protocol is used, creating a new path dependency. The cause and the effect are clear for protocol developers: ‘the incentive for QUIC was to try and prevent ossification in the network, but I think the implication is that it’s going to take power away from the network’²⁵. The reasons for this are the limitation incurred by the ossified network and power imbalance: ‘I do think that there’s a massive power differential that exists between people who run the network and the end users’²⁶, and now ‘the pendulum is swinging the opposite way’²⁷ back to the end users. While QUIC restores the end-to-end principle, it cannot overcome the differentiation between users and providers introduced through NAT directionality, and therefore it does not restore equality between all hosts.

This brings us back to the initial conception of the Internet architecture’s imaginary, wherein all hosts were equal and one could freely deploy protocols, strutted by architectural principles like end-to-end, permissionless innovation, and openness. While the limitations of ossification have partially been overcome through QUIC, this has only been possible by a significantly resourced transnational corporation, that also controlled large parts of the network, and controlled the worlds most-used browser, and could hire the best engineers, some of whom previously extensively worked on SCTP.

In other words, a precondition to restore part of the Internet architecture imaginary, was a significant economic incentive and technical and economic concentration, which contributes to a even further consolidated technological and economic reality. The increasing dominance of socioeconomic considerations over sociopolitical considerations is illustrated a member of the senior IETF leadership who confirmed that: ‘you need to play into some of the operators or vendors earning models in order to get something deployed’²⁸. This reflects demographic changes in the IETF: whereas the IETF used to be dominated by researchers, the overwhelming majority of participants now is representing the private sector. Deploying new protocols is still possible on higher layers of the networking stack, but less so in the lower layers of the architecture, unless one can gather resources like the one Google could muster, as we have seen in the example of SCTP. One simply needs to abide by the rules set by transnational corporations. In the words of a long time participant in the Internet protocol community: ‘[t]he mantra of the Internet enterprise is simple: ‘Get Big or Get Bought!’ (Huston, 2017: 5).

It is probably the same phenomenon we see in other industries. When it is brand new[...] you have more freedom to think exactly about how you want this thing to work and not worry about how much money you’re going to make, because you’re you will just make a lot of money. Now, a lot of people have made a lot of money off the Internet, there is still more, gobs and gobs need to be made, but it is a little bit crowded. We are heading a bit into the ‘winner-take-all’-phase.²⁹

While this startling socioeconomic reality, with significant impact on the material affordances of the Internet architecture, is widely recognized in interviews, the official IETF position is ‘to avoid policy and business questions, as much as possible’³⁰. This new socioeconomic reality, which is in part produced by the networks affordances, has cannibalized the Internets architectural imaginary, that is

²⁵ N2218

²⁶ N2218

²⁷ N2218

²⁸ N0118

²⁹ N0118

³⁰ <https://www.ietf.org/about/participate/get-started/> accessed on November 19, 2018

still being professed. The hollowed out sociotechnical imaginary actually functions as a cover, or even an implicit justification, for this consolidation of communication power.

Conclusion

The Internet architecture is hailed as an architecture in which all hosts are equal and everyone has the freedom to deploy their own protocols. This sociotechnical imaginary is anchored in the principles of end-to-end, permissionless innovation, and openness, and was operationalized through a process of co-production in Internet governance institutions such as IETF. When the Internet architecture was privatized a tussle over control over datastreams ensued between networks operators, enabled by equipment vendors, and content providers. This tussle led to the reconfiguration of the technical affordances of the Internet architecture. In this reconfiguration the equality of hosts and the ability to deploy new protocols between hosts has been subverted.

The sociotechnical Internet architecture imaginary and its self-regulatory governance model have not been able to cement the equality of Internet hosts and the freedom of researchers, small companies or individuals to functionally deploy new protocols, especially on lower layers of the protocol stack. Previously central sociopolitical conceptions and considerations that were part and parcel of the architecture's sociotechnical imaginary effectively faded into the background, while socioeconomic considerations have acquired a far more prominent place in the shaping of the Internet's technological affordances. The Internet architecture imaginary, that is still professed in IETF, obscures the socioeconomic reality in which interoperation between transnational corporations has overshadowed the practice and ethos of doing things for 'the good of the Internet' (Mathew, 2014). This dynamic arguably has contributed to the relative absence of the Internet architecture in current academic and policy debates on government regulation of the Internet, whereas the much discussed platforms and search engines are only a part of the Internet power assemblage.

The preceding analysis has shown how economic drivers spurred iterative changes in the affordances and materiality of the Internet architecture as well as its sociotechnical imaginary, illustrating the dynamic interrelation between economic drivers and technological affordances. This analysis contributes to the debates in governance studies by concluding that the self-regulation of the Internet architecture undermined the very design goals of the Internet architecture, changed its sociotechnical imaginary, and facilitated the prioritization of corporate interests.

Future research could focus on how standards development and self-regulatory governance bodies can take explicit values, such as the equality and freedoms of users, structurally into account as a formalized part of their processes, and what internal or external incentive structures would be needed to achieve this.

Funding Acknowledgment

This project has benefited from funding from the Ford Foundation (grant agreement 132156); the European Research Council under the European Union's Horizon 2020 research and innovation program (grant agreement 639379-DATACTIVE), and the Amsterdam School of Cultural Analysis Finishing Fellowship.

- Abbate J (1999) *Inventing the Internet*. Inside technology. Cambridge, MA: The MIT Press.
- Arpagian N (2016) The Delegation of Censorship to the Private Sector. In: Musiani F, Cogburn DL, DeNardis L, et al. (eds) *The Turn to Infrastructure in Internet Governance*. Information Technology and Global Governance. New York: Palgrave Macmillan US, pp. 155–165. DOI: 10.1057/9781137483591_8.
- Braden R (1989) *RFC1122 - Requirements for Internet Hosts - Communication Layers*. IETF. Available at: <https://tools.ietf.org/html/rfc1122> (accessed 30 July 2019).
- Braman S (2011) The Framing Years: Policy Fundamentals in the Internet Design Process, 1969–1979. *The Information Society* 27(5): 295–310. DOI: 10.1080/01972243.2011.607027.
- Braman S (2012a) Internationalization of the Internet by design: The first decade. *Global Media and Communication* 8(1): 27–45. DOI: 10.1177/1742766511434731.
- Braman S (2012b) Privacy by design: Networked computing, 1969–1979. *New Media & Society* 14(5): 798–814. DOI: 10.1177/1461444811426741.
- Braman S (2013) Laying the path: governance in early internet design. *info* 15(6): 63–83. DOI: 10.1108/info-07-2013-0043.
- Carpenter BE (2000) *RFC2775 - Internet Transparency*. Available at: <https://tools.ietf.org/html/rfc2775> (accessed 30 August 2018).
- Carpenter BE and Brim SW (2002) *RFC3234 - Middleboxes: Taxonomy and Issues*. Available at: <https://tools.ietf.org/html/rfc3234> (accessed 29 August 2018).
- Castells M (2009) *Communication Power*. Oxford, UK; New York: Oxford University Press. Available at: <http://public.eblib.com/choice/publicfullrecord.aspx?p=472226> (accessed 19 January 2016).
- Cath C and Floridi L (2017) The Design of the Internet’s Architecture by the Internet Engineering Task Force (IETF) and Human Rights. *Science and Engineering Ethics* 23(2): 449–468. DOI: 10.1007/s11948-016-9793-y.
- Chadwick A (2006) *Internet Politics: States, Citizens, and New Communication Technologies*. 1 edition. New York: Oxford University Press.
- Clark D (1992) A Cloudy Crystal Ball - Visions of the Future. In: *Proceedings of the Twenty-Fourth Internet Engineering Task Force*, July 1992, pp. 539–543. Available at: <http://www.ietf.org/proceedings/24.pdf>.
- Clark D (2018) *Designing an Internet*. 1 edition. Cambridge, MA: The MIT Press.
- Clark D, Wroclawski J, Sollins KR, et al. (2005) Tussle in Cyberspace: Defining Tomorrow’s Internet. *IEEE/ACM Transactions on Networking* 13(3): 462–475. DOI: 10.1109/TNET.2005.850224.

- Cowhey PF, Aronson JD and Richards J (2009) Shaping the architecture of the US information and communication technology architecture: A political economic analysis. *Review of Policy Research* 26(1–2): 105–125.
- Crowcroft J (2007) Net Neutrality: The Technical Side of the Debate: a White Paper. *SIGCOMM Comput. Commun. Rev.* 37(1): 49–56. DOI: 10.1145/1198255.1198263.
- Davidson A, Morris J and Courtney R (2002) Strangers in a strange land: Public interest advocacy and internet standards. In: *Telecommunications Policy Research Conference, Arlington, VA, September, 2002*.
- Davies DW, Bartlett KA, Scantlebury RA, et al. (1967) A Digital Communication Network for Computers Giving Rapid Response at Remote Terminals. In: *Proceedings of the First ACM Symposium on Operating System Principles*, New York, NY, USA, January 1967, pp. 2.1–2.17. SOSP '67. ACM. DOI: 10.1145/800001.811669.
- Deering SE (1989) *RFC1112 - Host extensions for IP multicasting*. Available at: <https://tools.ietf.org/html/rfc1112> (accessed 19 November 2018).
- DeNardis L (2009) *Protocol Politics: The Globalization of Internet Governance*. Cambridge, MA: MIT Press.
- DeNardis L (2014) *The Global War for Internet Governance*. New Haven: Yale University Press.
- DeNardis L and Musiani F (2014) *Governance by Infrastructure: Introduction, 'The Turn to Infrastructure in Internet Governance'*. ID 2730689, SSRN Scholarly Paper, 15 September. Rochester, NY: Social Science Research Network. Available at: <https://papers.ssrn.com/abstract=2730689> (accessed 27 February 2017).
- Denzin NK, Lincoln YS and Giardina MD (2006) Disciplining qualitative research. *International Journal of Qualitative Studies in Education* 19(6): 769–782. DOI: 10.1080/09518390600975990.
- Ding AY, Korhonen J, Savolainen T, et al. (2013) Bridging the gap between internet standardization and networking research. *ACM SIGCOMM Computer Communication Review* 44(1): 56–62. DOI: 10.1145/2567561.2567572.
- Dourish P (2018) The once and future internet: infrastructural tragedy and ambiguity in the case of IPv6. *Internet Histories* 2(1–2): 55–74. DOI: 10.1080/24701475.2018.1455311.
- Easterling K (2014) *Extrastatecraft: The Power of Infrastructure Space*. Verso Books.
- Eeten MJ van and Mueller M (2013) Where is the governance in Internet governance? *New Media & Society* 15(5): 720–736. DOI: 10.1177/1461444812462850.

- Epstein D, Katzenbach C and Musiani F (2016) Doing internet governance: practices, controversies, infrastructures, and institutions. *Internet Policy Review*. Available at: <https://policyreview.info/articles/analysis/doing-internet-governance-practices-controversies-infrastructures-and-0> (accessed 8 April 2018).
- Flanagin AJ, Flanagin C and Flanagin J (2009) Technical code and the social construction of the internet: *New Media & Society*. DOI: 10.1177/1461444809341391.
- Flichy P (2007) *The internet imaginaire*. Cambridge, Mass.: MIT Press.
- Francis P and Egevang K (1994) *RFC1631 - The IP Network Address Translator (NAT)*. Available at: <https://tools.ietf.org/html/rfc1631> (accessed 4 October 2018).
- Freed N (2000) *RFC2979 - Behavior of and Requirements for Internet Firewalls*. Available at: <https://tools.ietf.org/html/rfc2979> (accessed 30 August 2018).
- Galloway AR (2006) *Protocol: How Control Exists After Decentralization*. New Ed edition. Cambridge, Mass.: MIT Press.
- Hanseth O and Monteiro E (1997) Inscribing behaviour in information infrastructure standards. *Accounting, Management and Information Technologies* 7(4): 183–211. DOI: 10.1016/S0959-8022(97)00008-8.
- Haraway DJ (1988) Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. *Feminist Studies* 14(3): 575–599. DOI: 10.2307/3178066.
- Haraway DJ (1991) *Simians, Cyborgs, and Women: The Reinvention of Nature*. New York: Routledge.
- Huston G (2017) The Internet's Gilded Age. In: *The ISP Column*. Available at: <http://www.potaroo.net/ispcol/2017-03/gilding.html> (accessed 27 August 2018).
- Hutchby I (2001) Technologies, texts and affordances. *Sociology* 35(2): 441–456.
- Internet Architecture Board (1996) *RFC1958 - Architectural Principles of the Internet*. Available at: <https://tools.ietf.org/html/rfc1958> (accessed 29 August 2018).
- Internet Architecture Board (2004) *RFC3724 - The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture*. Available at: <https://tools.ietf.org/html/rfc3724> (accessed 29 August 2018).
- Internet Society (2012) Internet Invariants: What Really Matters. In: *Internet Society*. Available at: <https://www.internetsociety.org/internet-invariants-what-really-matters/> (accessed 26 June 2018).
- Internet Society (2013) The Value of Openness for a Sustainable Internet. In: *Internet Society*. Available at: <https://www.internetsociety.org/resources/doc/2013/the-value-of-openness-for-a-sustainable-internet/> (accessed 5 September 2018).

- Jasanoff S (2004) *Ordering Knowledge , Ordering Society*. In: Jasanoff S (ed.) *States of Knowledge: The Co-Production of Science and the Social Order*. London: Routledge.
- Jasanoff S and Kim S-H (2015) *Dreamscapes of Modernity: Sociotechnical Imaginaries and the Fabrication of Power*. University of Chicago Press.
- Lave J (2011) *Apprenticeship in Critical Ethnographic Practice*. University of Chicago Press.
- Lessig L (2006) *Code 2.0*. New York: Basic Books. Available at: <http://codev2.cc/download+remix/Lessig-Codev2.pdf>.
- Levinson NS and Cogburn DL (2016) The Next “Turn” in Internet Infrastructure Governance. In: Musiani F, Cogburn DL, DeNardis L, et al. (eds) *The Turn to Infrastructure in Internet Governance*. Information Technology and Global Governance. New York: Palgrave Macmillan US, pp. 219–223. DOI: 10.1057/9781137483591_12.
- Mansell R (2013) *Imagining the Internet: Communication, Innovation, and Governance*. Oxford: Oxford University Press.
- Mathew AJ (2014) *Where in the World Is the Internet? Locating Political Power in Internet Infrastructure*. Berkeley: University of California. Available at: <https://www.ischool.berkeley.edu/research/publications/2014/where-world-internet-locating-political-power-internet-infrastructure> (accessed 30 January 2018).
- McKelvey F (2018) *Internet Daemons: Digital Communications Possessed*. U of Minnesota Press.
- McLuhan M (2013) *Understanding Media: The Extensions of Man*. Critical edition. Gingko Press.
- Minar N and Hedlund M (2001) A network of peers: Peer-to-peer models through the history of the Internet. In: *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*. O’Reilly Media, pp. 9–20.
- Mogul JC, Theimer M, Finlayson R, et al. (1984) *RFC903 - A Reverse Address Resolution Protocol*. Available at: <https://tools.ietf.org/html/rfc903> (accessed 19 November 2018).
- Mosco V (2005) *The Digital Sublime: Myth, Power, and Cyberspace*. The MIT Press.
- Musiani F, Cogburn DL, DeNardis L, et al. (eds) (2016) *The Turn to Infrastructure in Internet Governance*. New York, NY: Palgrave Macmillan.
- Orwat C and Bless R (2016) Values and networks: Steps toward exploring their relationships. *ACM SIGCOMM Computer Communication Review* 46(2): 25–31.
- Padlipsky MA (1982) *RFC871 - Perspective on the ARPANET reference model*. IETF. Available at: <https://tools.ietf.org/html/rfc871> (accessed 30 July 2019).

- Price M and Verhulst S (2000) The Concept of Self-Regulation and the Internet. *Protecting Our Children on the Internet: Towards a New Culture of Responsibility*: 133–198.
- Resnick P (2014) *RFC7282 - On Consensus and Humming in the IETF*. IETF. Available at: <https://tools.ietf.org/html/rfc7282> (accessed 30 July 2019).
- Rogers M and Eden G (2017) Digital Citizenship and Surveillance| The Snowden Disclosures, Technical Standards, and the Making of Surveillance Infrastructures. *International Journal of Communication* 11(0): 22.
- Russell AL (2006) 'Rough Consensus and Running Code' and the Internet-OSI Standards War. *IEEE Annals of the History of Computing* 28(3): 48–61.
- Russell AL (2014) Open Standards and the Digital Age: History, Ideology, and Networks. *Afterimage* 42(3): 39.
- Rüth J, Poese I, Dietzel C, et al. (2018) A First Look at QUIC in the Wild. *arXiv:1801.05168 [cs]*. Available at: <http://arxiv.org/abs/1801.05168> (accessed 31 August 2018).
- Saltzer J, Reed D and Clark D (1984) End-to-end arguments in system design. *ACM Transactions on Computer Systems (TOCS)* 2(4): 277–288. DOI: 10.1145/357401.357402.
- Shapiro C and Varian HR (1998) *Information Rules: A Strategic Guide to the Network Economy*. Boston, Mass: Harvard Business Review Press.
- Smyrniaios N (2018) *Internet Oligopoly: The Corporate Takeover of Our Digital World*. Emerald Publishing Ltd.
- Taylor T, Schwarzbauer HJ, Kalla M, et al. (2000) *RFC2960 - Stream Control Transmission Protocol*. Available at: <https://tools.ietf.org/html/rfc2960> (accessed 31 August 2018).
- ten Oever N (2019) Productive Contestation, Civil Society, and Global Governance: Human Rights as a Boundary Object in ICANN. *Policy & Internet* 11(1): 37–60.
- ten Oever N and Cath C (2017) RFC8280 - Research into Human Rights Protocol Considerations. *IRTF*. DOI: DOI 10.17487/RFC8280.
- Thaler D (2011) *RFC6250 - Evolution of the IP Model*. Available at: <https://tools.ietf.org/html/rfc6250> (accessed 2 October 2018).
- Van Schewick B (2012) *Internet Architecture and Innovation*. Mit Press.
- Zittrain J (2008) *The Future of the Internet—and How to Stop It*. Yale University Press.