‘This is not how we imagined it’
Technological Affordances, Economic Drivers and the Internet Architecture Imaginary

Abstract
While Internet platforms have come under public scrutiny due to privacy invasion and the spread of misinformation, the Internet architecture is still generally perceived as an engine for innovation and as an enabler of human rights. This is based on a prevailing socio-technical imaginary that hails back to the beginning of the ’90s and that is still prominent in one of the main Internet technical governance bodies, the Internet Engineering Taskforce (IETF). In this paper, I answer the question how technological and economic developments have impacted the Internet architecture imaginary by drawing on Science and Technology Studies and International Political Economy, using a mixed method approach based on document analysis, interviews and participant observation. Firstly I show how the Internet architecture is co-produced through a socio-technical imaginary based on notions such as: the end-to-end principle, permissionless innovation, openness, and the discussion of rights and liberties since its early beginnings. Secondly I show how the introduction of network management devices by network operators and network equipment vendors increased their control over data streams, violated the end-to-end principle, and altered the architectures affordances and hampered innovation. Thirdly I explore the iterative attempts by content providers to innovate on the transport layer, reinstate the end-to-end principle and regain control over data-flows. I make two main observations: firstly, both swings of the architectural and power pendulum on the transport layer, from endpoints, to network operators and back, coincide with a significant concentration of market power and challenges to the architecture imaginary. Secondly, with the new technical and economic ordering of the Internet architecture, the Internet protocol community arrived at a crossroads; will it either restore the Internet architecture imaginary, or go further down the path of economic realism.

Introduction
When in the early nineties the Internet escaped the labs and found its way to millions of users, it was widely perceived as an engine of innovation (Van Schewick 2012), an information highway (Flichy 2007) and a tool for democratization (Castells 2009). These economic, social and political aspirations were rooted on a firm belief in the underlying architecture of the Internet, which would route around censorship, promote human rights (Carr 2013), and leverage all kinds of applications through its generative nature (Zittrain 2008). Here I seek to show how the Internet architecture imaginary, the combination of values, institutions and visions, developed after three decades of large scale use and significant technical and economic changes.

This paper interrogates the socio-technical architectural imaginary which is pervasive among Internet protocol engineers (Wyatt 2004; Flichy 2007). A socio-technical imaginary describes 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).

The Internet architecture is co-produced (Jasanoff 2004) by multi-national corporations, states, scientists and advocates in a self-regulating governance body called the Internet Engineering Taskforce (IETF) (Poullet 2001). 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 development of the architecture of the Internet has taken place through the development of open and voluntary standards to facilitate interoperability and compatibility between the products and implementations 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 do determine vectors of control over user data flows (Galloway 2006), how users access information (DeNardis 2014), and how users can exercise their rights online (Lessig 2006). This begs the questions **how the Internet architecture and its imaginary**

adapts to technical and economic changes since its commercialization. I seek to answer this question by tracing how the socio-technical architectural imaginary influenced the co-production of the Internet architecture through the IETF as self-regulated governance institution, the technology of the Internet architecture, the documentation of Internet standards, mailinglist conversations and meeting recording and the engineers that are engaged in the development of a rapidly evolving technical architecture. In this paper, I aim to make three contributions to Science and Technology Studies and Governance Studies literature. Firstly I contribute to the understanding of the evolution of the socio-technical imaginary of the Internet architecture through which the Internet architecture is co-produced by describing its development and reactions to a changing environment. Secondly, I contribute to the debate on governance through infrastructure (Musiani et al. 2015) by providing insight into how the materiality and affordances of the Internet architecture are intertwined with the power struggles over data streams among network operators and network equipment on the one hand and content providers on the other hand. Thirdly, I contribute to the debate at the intersection of governance studies and political economy on the ability of self-regulated sectors to act in the public interest (Den Hertog 2010).

I will start out by outlining the Internet architecture imaginary, and its main the values, visions and principles. Subsequently I will describe the first challenge to the architecture imaginary, the introduction of network management devices by network operators and how they changed the affordances of the architecture. Then I will continue to describe iterative responses to this

Methods

The object of my research is the Internet architecture, and the Internet infrastructure which it describes and produces. How can one study an architecture without falling to high level abstractions or getting lost in the details of the hundreds of thousands of interconnected networks that make up the Internet? Star and Ruhleder (1994) describe infrastructures as relations, which is particularly true for the Internet which an internetwork between autonomous networks. To understand the characteristics of this network I studied the community that produces the protocols for interoperation between these networks, which are discussed on mailinglists and during meetings, and finally published in the Request for Comments-series (RFCs). The RFC-series should not only be understood as a series of technical documents, but also as policy documents (Kahin, Keller, and Harvard Information Infrastructure Project 1995; Braman 2013) and therefore relevant for understanding the sociotechnical imaginary of the Internet architecture.

While engaging in participant observation for four years I gained deeper insight in how engineers worked as craftpersons (Sennett 2008; Losey and Meinrath 2016), negotiators (Singh 2008) and digital citizens (Isin and Ruppert 2015). Next to that I also got exposed to the workings of procedures and processes and the institutional ordering it produced, as well as the to the prevalent underlying economic interests.

I became involved in Internet protocol community during my role as Head of Digital for ARTICLE 19, an international not-for-profit freedom of expression organization. The work on this article and theory building started as an ethnographic memoir while engaged in Internet Engineering Taskforce (IETF) and its sister organization the Internet Research Taskforce (IRTF) between March 2014 and July 2018. In this period I attended and participated in 11 IETF and IRTF meetings; IETF meetings are held three times per year for six days. In addition negotiation to meetings, discussions and coordination takes place on mailing lists. I was involved as a civil society actor within the IETF and IRTF where I actively participated in several Working Groups and Research Groups, served as co-chair of the Human Rights Protocol Considerations Research Group, member of the Internet Research

Steering Group, editor of the 'Tao of the IETF', and as member of the Education, Mentoring and Outreach Directorate.

This experience provided me with a first-hand account of the practices in the most prominent Internet standards and research body, as well as access that an external observer might not otherwise have. On the other hand, my direct involvement might also have introduced some biases vis-à-vis the IETF and IRTF and actors engaged in it. To counter this potential bias, I reflected on my own role in the process, and by inferring, checking and validating insights through mixed methods and a variety of sources.

In order to gain insight in the Internet architecture and how it is produced, I engaged in quantitative analysis of the RFC series and IETF mailinglists focusing on prevalence and development over time of language related to public interest, ethics and rights and the trends in the professional affiliation of document authors. I obtained these documents from the IETF website, where these documents were publicly available at the time of writing of this article. I undertook a quantitative mailing list analysis using the Python-based tool BigBang,[1] to gain a deeper insight into trends and interactions, which formed the basis for the questionnaire I developed for the interviews. I engaged in qualitative analysis through participant observation during 11 IETF meeting and on mailinglists. 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 interviews took place during IETF meetings between March and July 2018. I refer to interviewees position and role in the IETF and IRTF when it is relevant for the argument, and when the interviewees provided written consent for doing so. The default position was to disguise the personal identities of research participants, unless the interviewees consented in writing (Saunders, Kitzinger, and Kitzinger 2015). Both audio and transcriptions were encrypted, stored and processed with free and open source software. The audio was transcribed and analyzed using qualitative methods informed by thematic analysis (Guest, MacQueen, and Namey 2012), 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 developed from the arly '90s up to now and what the drivers, vision and struggles were.

To validate and triangulate the findings from the quantitative mailinglists analysis, RFC analysis and interviews, I engaged in the document analysis of RFCs. I made a purposive sample of 20 RFCs that mention 'architecture', and made a chain sample 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 also analyzed 5 RFCs that got mentioned during the interviews.

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 burgler to distract the watchdog of the mind' (McLuhan 2013, 19) I argue we should rather look at the design of the infrastructure which determines the shape of data stream; the Internet architecture.
The Internet architecture consists of 'standards which make up the technical back-bone of an information infrastructure' (Hanseth and Monteiro 1997, 183). The Internet architecture, through its affordances, possibilities that it allows for (Gibson 1978), shapes our information societies, but it is also still being changed and developed. The Internet architecture is co-produced between multi-national corporations, states, scientists and advocates in governance bodies and Standards Developing Organizations (SDOs) such as the Internet Engineering Taskforce. This process of co-production entails

[1] http://datactive.github.io/bigbang/ accessed on August 25, 2018

that visions, values and technological artifacts are developed together with institutions, identities and discourses (Jasanoff 2004). The normative aspects of this collective imagination combined with the institutions and the materiality of socio-technical endeavors can be understood as its socio-technical imaginary (Jasanoff and Kim 2015). This means that to understand the impact of the Internet architecture on our societies, it is important to understand the process of co-production, and the sociotechnical imaginary that guides it. 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), but at the same time the prevalent socio-technical imaginary nurtures a collective and individual identity, and is anchored in documents, procedures of the IETF (Mosco 2005; Kelty 2005). The imaginary revolves around doing things that are 'for the good of the Internet' (Mathew 2014, 160), sustaining a 'generative Internet' (Zittrain 2008, 6), having a 'nostalgia for the future' (Carey 1989, 200), and revolves around specific engineering principles (Internet Society 2012). This idealistic representation starkly contrast with how DeNardis (2009) describes how the mere transition from one protocol to another can cause a significant amount of contestation because of its geopolitical and financial interests and impact. To add to the complexity, Braman (2009) shows how discussions about rights, freedoms and social policy have been part of the Internet standards process since it beginning. Calls for the formal integration of rights consideration into the standards process have never ceased (Morris and Davidson 2003), but are experiencing a recent uptick due to both the Snowden revelations and the increasing impact of the Internet on society (Cath and Floridi 2017; Rachovitsa 2017). This has also lead to a debate about the legitimacy and adequacy of the self-governing technical standards bodies (Rogers and Eden 2017).
Both in the academic literature as well as in the standards body itself there is a discussion about economic consolidation. The commercialization of the Internet in the '90s together with its technical and architectural affordances such as its modularity and layering principles, was expected to lead to perfect markets, free competition and decentralized structures (Litan and Rivlin 2001). This did not happen. Rather, 'market concentrations, control and power struggles are categories to adequately describe the fundamental dynamics of the commercial Internet' (Dolata and Schrape 2018, 85). While the commercialziation was supposed to lead to competition and innovation (Cowhey, Aronson, and Richards 2009; Van Schewick 2012), it actually led to emerging Internet oligopolies (Mansell and Javary 2002). What this means for the Internet on the platform and application level has been studied to a great extend both in terms of political economy, as well as the technological materiality, but this is not the case for the Internet architecture, which could be one of the reasons why the Internet architecture imaginary is still prevalent. By combining the analytical lenses of International Political Economy (IPE) and Science and Technology Studies (STS) I aim to foreground the effect of self-regulation on the Internet architecture imaginary. The combination allows me to take into account economic drivers, the materiality of the technology, without underemphasizing either of them.

Analysis - On the idea of smart endpoints and the dumb pipes
In this analysis I will first describe the Internet architecture imaginary and trace its evolution in the light of the following developments; the rise of the middlebox and the changing affordances of the architecture, the failed deployment of the Stream Control Transmission Protocol (SCTP), and the successful development and deployment of the Quick UDP Internet Connections (QUIC) protocol.

The Internet architecture imaginary
The socio-technical Internet architecture imaginary emerged during the early phases of the development of the Internet while it was still a research network. Here I focus on the development of stabilization that came after that during the commercialization of the Internet, when the US government

ceased control and gave way for an increased amount of self-regulation through private governance bodies such as the IETF. Here I will describe some of the characteristics of the Internet architecture imaginary which helped co-produce the actual Internet architecture, its institutions, and technology. When asking engineers in the IETF for the central architectural values or principles of the Internet protocol community, their answers have a surprising amount of overlap. The end-to-end principle (Saltzer, Reed, and Clark 1984), openness, and permissionless innovation get mentioned time and again in interviews as well as technical and policy documents (Internet Society 2012).

In almost all interviews the end-to-end principle came back as a pillar of the values and principles of the architecture, 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 (Carpenter 1996), and let the network be 'dumb pipes'[2] that only transport packets. The importance of the principle, or argument, was understood as that it 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, um, from anywhere to anywhere.'[3] This principle was infra-structurally also a revolution, especially because it contrasted so strongly with the telephony networks that preceded the Internet; that the exact opposite of the end-to-end principle, namely a smart network with little functionality in end devices, also known as 'dumb edges'. The ablility and design to have smart edges in a network with only one functionality, namely transporting packets, would also allow for a lot of tinkering experimentation, which gave birth to the identity digital craftsperson (Losey and Meinrath 2016). Changes at endpoints that are under complete control of their owners can be changed quickly, both in terms of hardware as well as software. Changes in the infrastructure would be a large investment, or in the words of a former telecoms 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'[4]. This captures the importance of the end-to-end principle for innovation, but it also has further implications as a socio-technical conceptualization, one interviewee 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.'[5]

This gave rice to this community of tinkerers who had the ability to shape and create their own computing and networking experience, be it alone, in groups, within their own institutions or in their company. This had a tremendously empowering effect on engineers, because their main impediment was no longer access to a large mainframe or expensive computing time, but rather what engineers could let computers do. 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'[6].

The principle of permissionless innovation is closely tied to the end-to-end principle. Where the end-to-end principle talks about where functionality should be placed in the network, the permission innovation principle describes that there are no barrier for the deployment of new protocols. In other words: 'you should have the functionality that you need as an endpoint so you don't need to negotiate with any entity in the middle of the network to get your new thing deployed. The interfaces to the

[2] N1418
[3] N0218
[4] N0918
[5] N0618
[6] N1518

underlying layers are sufficiently specified that you build an application on top of it you don't need to negotiate with any entity in the middle of the network in order to transport your packets.'[7] Negotiating here is meant here in both technical terms as well as in political terms, permissionless innovation made it clear that there should be no governance body which can sanction what protocols should be used. This echoes both the history of the telecommunications networks, as well as the tries of the International Telecommunication Union to gain control over the nework as well as protocol development. Large applications on the Internet are often used as typical examples of permissionless innovation, for instance: 'a typical example is the web. Tim Berners-Lee did not ask permission from anyone, he invented some thing, went back, built it, and then it was downloaded and no one at the IETF or ICANN[8] or government had anything to say about it. After the web it was the same thing for the two important things: bittorrent and bitcoin.'[9] This socio-technical conceptualization again has clear socio-political implications; there should be no limitations on the ability to design, develop, deploy and use new technologies on the network. This principle extended also beyond the protocols that got developed in the IETF, the community members are very clear on that point: 'we're not protocol police'[10].

Openness was mentioned as the principle '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'[11], which was elaborated on by another IETF community member in a leadership position:

> 'The first principle that comes to mind is openness. Um, that's something that hasn't changed since the beginning that once we started establishing this network, the idea would be that it would be easy to put a new node on the network. It would be easy for someone to come in and develop software, hardware, applications for that and that any attempt to shut down parts of the Internet were easy to route around and keep everything connected and all of that as far as I can tell, it's still there and we see that when we get things like the Arab spring stuff where Egypt shut down the internet in Egypt and they were able to do that in their isolated domain, but they couldn't shut it down outside. It didn't affect anybody outside their realm.'[12]

While there is a rather clear socio-technical conception, here it is directly coupled again with a socio-political imaginary of connectivity, access, as well as the explicit political consequences. Except for the openness of the network, openness is also often associated with the consensus approach to standards development, which fits into an 'ideology of open standards' (Russell 2014, 21) which '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, as the Internet Society lays out in their paper 'The Value of Openness for a Sustainable Internet' (2013). The ideal for participatory democracy and open governance is also reflected in the IETF's unofficial motto, coined by Dave Clarke: *'we reject: kinds and presidents, we believe in: rough consensus and running code'* (Ten Oever and Moriarty 2018).This shows the strong interrelation between the technology, the institution and community that co-produce the architecture. But it also shows how many terms have both a technical and a political connotation; another example is 'transparency', which often gets mentioned in conjunction with openness. In relation to standards and

[7] N0118
[8] Internet Corporation for Assigned Names and Numbers, the Internet governance and coordination body for the Domain Name System and the allocation of IP addresses
[9] N1618
[10] N0218
[11] N0918
[12] N0418

governance relates transparency related to transparent methods for rule-making in the standards process, for instance through open mailinglists archives and video archives of meetings, but in technical terms Internet transparency refers to a 'network that does not filter or transform the data that it carries' (Aboba, Davies, and Davies 2007, 1). Another example is the term connectivity, RFC1958 goes as far as to describe that:

> 'Many members of the Internet community would argue that there is no architecture, but only a tradition […] the community believes that the goal is connectivity, the tool is the Internet Protocol, and the intelligence is end to end, rather than hidden in the network.' (Carpenter 1996).'

The Internet protocol community describes the network both a technical artifact, as well as something that escapes it, in which 'connectivity is its own reward' (Carpenter 1996, 2). A metaphor that is also often used to describe this vast area of protocols, networks, software, hardware, vendors, operators and providers, is the term *ecosystem* (Thomas and Wyatt 1999; Wyatt 2004). Organic metaphors such as 'ecosystem'[13], 'healthy architecture'[14], 'ossification'[15] invoke an almost cosmic inevitability of the Internet, but at the same time call for a care for it. When brought up someone remarked that organic metaphors are used because 'no one fully understands the whole netfriworking stack, similar to how we do and don't understand nature at the same time'[16]. What this adds to the imaginary is that next to being described as being complex, the architecture is also described as constantly changing (Carpenter 1996; Braman 2016). This seems to imply that the architecture changes by itself, but rather it are engineers that are changing the networks and the protocols. Almost all network engineers refer back to RFCs or are even part of the IETF community by partaking in IETF mailinglists. Analysis of mailinglists show that on average people who send an email to an IETF mailinglist stay active for roughly five years, but many stay on for longer. Active participants often participate in the physical meetings, which take six days are are held three times per year. This creates a quite tight knit community in which individuals formally represent their own opinions and not those of their employer, which add to the sense that once you participate, you are part of an institution, a project, that is larger than you or your job.

In this first part I have sketched different part of the Internet architectural imaginary and the community that still shapes it and gets shaped by it. We have seen how technical concepts such as the end-to-end principle, openness and permissionless closely interrelate with political implications and conceptions as well as economic realities. In the next part I will show how the imaginary withered technical and economic changes.

Cracks in the imaginary: The Rise of the Middlebox

The first threat to the end-to-end principle and the openness of the Internet took the form of '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), the shorthand for these devices became *middleboxes.* Middleboxes can have many different functions. Common uses of middleboxes are firewalls, NAT routers, IP-tunnels (such as VPNs), caching servers and network management devices. With the introduction of middleboxes there came a paradigmatic shift in the functioning of the network. Whereas the network was previously supposed to function as a

[13] N0318, N0618, N1018, N1118
[14] N1118
[15] N1118, N1018, N1318, N1418, N1718
[16] Fieldwork notes March 2017

'dumb pipe'[17], as outlined in the end-to-end principle (Saltzer, Reed, and Clark 1984; Carpenter 1996; Bush and Meyer 2002), now functionality was added to the network. This happened because of three effects of the rapid growth of the network: (1) increased need for security, (2) the depletion of 32 bits IPv4 addresses, and (3) increasing business interests (Internet Architecture Board 2004). These developments cannot be reduced to either technical or economic affordances; if the Internet would have been designed with security as one of it design priorities, it might have not lead to the need for middleboxes. But having security built in might have hampered other experiments or developments that would have increased barriers for the openness of the early Internet. Same holds true for the 32 bits IPv4 address, one could easily think that the early network designers should have simply picked a larger number, but this would have hampered the amount of information that could flow through the network on relatively low bandwidths, if every packet needed to have a large address in its header. This would have made the Internet less usable, with might have contributed to a slower uptake, and thus larger addresses and middleboxes would not have been needed, and extra business interests would not have been generated. So there are clear path dependencies as well as interdependencies that led to the growth of the network.

With the growth of the network, it changed from a research network to a network that was connected to critical services that could not go offline. With the network growing beyond a relatively trusted environment of a group of trusted 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 by at home routers, but also inside larger networks and were thus not only found on the edges of the Internet. A regularly implemented functionality of firewalls is directionality. This meant that a network would be 'protected' from receiving connections from an outside computer that it did not explicitly ask for. This sounds like a sound security measure on the one hand, but what it does it that it makes a difference between clients and servers. If you behind a directional NAT, you cannot function like a server because you need to be the device that initiates the connection. While many smartphones currently have more processor capacity and storage space than early webservers, they cannot function as one simply because the network is imposing another ordering. This is how firewalls differentiate between produces and consumers and creates users. Another aspect of middleboxes is that they are not detectable or advertise themselves and their functionality, so they change the affordance of the network, without declaring why, how or on whose behalf.

Another consequence of the growth of the network was that the unique numbers for hosts connected to the network was running out since it was never envisaged that so many devices would be connected to the network. There was no direct replacement addressing scheme ready (later this would become IPv6 (DeNardis 2009)), and there were projections that IPv4 addresses would run out in 1994[18]. This lead to the introduction of Network Address Translation which created a number of addresses for used inside a network, which there is only one IP address for the Internet (Srisuresh and Holdrege 1999). This specifically went against the Internet transparency principle according to which in 'packets flow unaltered through the network' (Carpenter 2000), but it did address a concrete technical problem. Because this meant that for instance you would not need an IP address for all devices in your house, but your home router would be the only device with a publicly facing IP address, and the home router would assign addresses to devices within the home network. These NATs were not only implemented at edge networks, such as home routers, but also in intermediate networks, called carrier-grade NATs. This makes the packet flow through the Internet less transparent because the IP address of the end-device is not known.

[17] N1018
[18] N0718

A final reasons for the introduction of middleboxes was network management for the prioritization of services. On the one hand this could be used for so called ‘quality of service’ management, on the other hand this also turned to be a way for Internet Service Providers (ISPs) to differentiate services and prioritize some services over others, or block other services. A discussion more commonly know as the *net neutrality debate* (Crowcroft 2007). While it might seem efficient to prioritize video traffic over e-mail, since you probably would not notice when an email arrives a few second later,

The introduction of the middleboxes solved some immediate problems such as security issues, delayed other problems, such as the IPv4 shortage, and create some economic incentives, in the case of the prioritization of services, but it also created a direct barrier to innovation. Because middleboxes performed more functions than IP routers were initially meant to do, they filter traffic based on known characteristics. This directly impacted traffic flow and limited possibilities on the network because middleboxes refuse traffic that they do not recognize. This might be useful if it is the intended use, but in case of the developments of new protocols, one would want to have access to these, and in case the middlebox does not get updated, it is hard to find out who to contact.

Cracks in the imaginary; The Advent of Ossification

The Internet is defined by two main characteristic: packet switching and the Internet Protocol (IP). IP provides all connected devices with a number, and packet switching allow data to be cut up in smaller packets with a header with sufficient information to route the packet over the network. One of the first transport protocols on the Internet is know as the Transmission Control Protocol (TCP), TCP is very reliable, but it is not the most useful when a low latency is important. This is why the Stream Control Transmission Protocol was initially developed. It was used within telecom providers to adapt their own networks to the Internet Protocol, instead of over their own proprietary technologies. Within the telephony providers it worked, it is currently still present in every phone switch in the world. Because it was developed after TCP it contained a lot of improvements. For instance it can handle multiple streams of data at the same time, it could connect to multiple hosts, as well as some other features. SCTP was first standardized for telephone networks in 2000 (T. Taylor et al. 2000), it was adapted to be a general purposed Internet protocol in 2004 (Stewart 2007), received several updates, got a new version in 2007 (Stewart 2007) got several updates and received its the last updated in 2013 (Tuexen, Ruengeler, and Stewart 2013). Significant amounts of work has gone in the development of this protocol, unfortunately it never significantly worked on the Internet. The reality was that SCTP worked perfectly in the lab and lived up to all the design expectations, but it would not work *in the wild*, on the actual Internet because of ‘intelligent’ middleboxes had led to an inflexibility on the Internet which is also known as *ossification.* NAT middleboxes were used to expect TCP traffic and did not recognize the the SCTP headers and would not let the SCTP traffic through (Hayes, But, and Armitage 2008). Even if the client initiated a request for SCTP traffic itself.

This shows how principles that lie at the core of the Internet imaginary cannibalized themselves; by allowing permissionless innovation on the network itself, the end-to-end principle got undercut, which in turn cannibalized on permissionless innovation.

Apart from hampering permissionless innovation, it also lead to new path dependencies. In the words of one interviewee:

> ‘firewalls didn't serve only a security purpose, they also served an administrative control purpose, that's a third party in the midst of the peers who are talking to each other. So it's been difficult for Internet peer to peer things to take off. There's plenty of peer-to-peer things that do take off through cleverness, like, you know, um, uh, there was an interesting moment when we discovered that bittorrent was one of the greatest user discoverers and understanders of capacity and that was not done here at the IETF, but people studied it and learned a lot from it and, but, but it had to sneak by various controls in order to do what it did. And um, and that type of protocols certainly exists, but we don't, we haven't engineered it so often because people like to administrative controls and vendors like to sell things that provide administrative controls.’ [19]

The introduction of middleboxes led to changed affordances in the network, with changed control from the end-point to the network operators, who were supported in this endeavor by networking equipment providers. This meant that agency was taken from the endpoint, and was added to the network. As I have shown this had both technical and economic reasons, which jointly surmounted to a reordering of the affordances of network.

The return of the endpoints

The control of the network operators who installed the middleboxes and the network equipment vendors who produced the middleboxes resulted in quite some resentment among the Internet protocol community, especially because it limited others in what they could do on the network. What added to the frustration that middleboxes also made it harder to troubleshoot issues since middleboxes do not announce themselves and hamper Internet transparency. But there were no protocol developer who could find a solution around the problem. SCTP developers had worked on the problem for almost a decade and did not solve it. For content providers, and especially rich multimedia content providers, it became increasingly important to have a protocol that would deliver more content in a faster manner over the network, because of the increasing demand in streaming video. This was only possible with the development of the Quick UDP Internet Connections (QUIC) protocol, a connection based stream protocol, which supported multiple streams, in similar to SCTP, coupled with some of the characteristic of TCP such as congestion control. Another quintessential difference between SCTP and QUIC was that the latter was developed by Google. Google already had a fast global Content Distribution Network (CDN) and developed and disseminated one of the most used browsers in the world; Chrome. In the words of an interviewee: ‘the reason that QUIC [..] can do what they can do is because the two end points are controlled by the same people, so they can, they can do like dark releases and AB-testing and all that that we can't do’[20]

The missing link for Google was a protocol to transport their content as fast and responsive as possible to their browsers. At the same time Google also wanted to be able to keep innovating with the protocol and regain control of the data flow, to ensure no competitors would have access to the data Google gather. To achieve this Google encrypted as much as possible, except for the necessary data in the packet header, thus keeping the content of the data and a lot of meta data hidden for the network. Previously the network has used this information to ‘manage’ the packet, prioritize services or do other things, now it would be much harder for network operators to understand what traffic is flowing over

[19] N0218
[20] N0218

their network. In this way network operators, or equipment vendors had a much harder time injection information, changing information or understanding what is actually inside the packet.
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 which using its 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. Instead of keeping QUIC for its own use, Google brought QUIC to the IETF for standardization, which would increase the chances of broader adoption, and therefore ensure that new middleboxes would not block QUIC traffic.

The development of QUIC is arguably going to lead to a new reordering of the network. Because it is possible to to penetrate middleboxes, and encrypts the traffic, network operators will have relatively little grip on the data and can harder shape or filter it. At the same time QUIC will also also for better connections on low bandwidth networks. Privacy concerns are also covered with the encryption that is on by default. There are a few things that QUIC cannot address, middlebox induced directionality is cannot be circumvented by QUIC because the connection still needs to be initiated by the endpoint. This reality on the network has created a differentiation between users and providers of data. This means that while the end-to-end principle gets restored, not all ends are created equal. The providers end are heavily consolidated in a few content and hosting providers, such as Google, Amazon, Facebook, Netflix. Ensuring a faster and more private and secure connection between these ends, will not necessarily lead to more innovation by the end user, even running a QUIC server stack is currently still very complicated.

So while the affordances of ossification might have partially been overcome in the case of QUIC, this has only been possible by a significantly resourced company such as Google, that controlled large parts of the network, provided a large browser and could hire the best engineers. So in overcoming this affordance there was a significant economic aspect that led to the reordering, which in turn also contributed to a new, and consolidated, economic reality.

Imaginaries They Are A Changin'

With the stabilizing of the Internet architecture imaginary at the start of its commercialization in the '90s, the rise of the middlebox in the '00s, the development of a response in the 10s, the questions begs what will happens in the '20s. While we consider the future of the architecture and its imaginary we should consider that '[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). If we view through these lens of economic reductionism, it seems that there is less space for individuals and smaller players. In the words of one interviewee: 'clearly permissionless innovation, the ability to to have a tremendously successful capability on the Internet is still real, even though now it's more likely to be over HTTPS rather than creating an entirely new protocol.' This means that innovation is still possible on the application layer, but less so in the lower layers of the architecture, unless resources like the one Google could muster need to be invested, as we have seen in the example of SCTP. And this trend towards consolidation can also be observed in the IETF, where the number of participating companies is dwindling and the companies are growing. In the words of one interviewee: '[t]he mantra of the Internet enterprise is simple: 'Get Big or Get Bought!''(Huston 2017, 5). And with the stabilizing of the Internet architecture in general, there might be less opportunities for growth;

> ‘It is probably the same phenomenon we see in other industries. When it is brand new, when there is green fields, and with a small amount of engineering effort you can create whole new markets and invent new products and market sectors that didn't exist before. Then 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. I think it is more of a challenge to differentiate your idea is gonna be profitable, whereas previously you didn't need to worry about it too much.’[21]

At the same time the Internet imaginary is still prevalent, and the engineers at the IETF still have the zeal of craftpersons and engineering pride. But at the same time the community is aging, how will it adapt to these new realities?

Some of the problems that have created middleboxes in the first place are now being solved; IPv6 is finally being deployed, which means there are sufficient IP addresses to provide all devices in the world with a number, so NATs won’t be needed anymore, when IPv4 phases out. What will happen with the network after that remains to be seen; will the applications and devices be able to handle their own security, will the IETF incorporate human rights considerations? Currently there seems little appetite for this;

> ‘if I am proposing something that has a privacy win, he wouldn't oppose it because he is against privacy, he would just say: you gotta give me something else. Give me some efficiency win, some reliability win, some robustness win, some other thing that I do care about in order for me to be your ally.’[22]

Or in the words of the current IETF chair:

> *‘you need to play into the earning models [...] to get something deployed’*.

Conclusion

The Internet architecture imaginary revolves around openness and free flow of information, which would create an environment for a perfect market for innovation and rights to flourish together. The Internet governance regime, of which Internet standards setting is a part, is often understood as a governance innovation, and even as a safeguard for the public interest (Johnson, Crawford, and Palfrey 2004; Castells 2011). To shape this, self-regulation has been hailed as best approach to support the open and distributed network architecture of the Internet, in part because it allows swift responses to new technologies and architectural evolutions (Price and Verhulst 2000).

What we see in the case of the Request for Comment series and the Internet architecture is that architectural principles, rights and freedoms have significantly been discussed and documented, but when technical problems arose, quick technical fixes were developed and implemented by companies that did not serve the Internet architecture nor followed the Internet architecture imaginary. Because the IETF insists on not being ‘the protocol police’ and ‘permissionless innovation’, the imaginary was not

[21] N0118
[22] N0618

enforced and thus innovation on the lower layers in the Internet architecture now only seems possible by large consolidated parties. QUIC was able to succeed where SCTP failed because of significant investments and control over servers and a dominant browser, which is not something readily available for individuals, research institutions, and only for very large transnational corporations.

In the IETF this can be observed by companies getting larger and consolidation taking place over several layers of the stack and in different parts of the sector. Ranging from hosting to networking hardware, from CDNs to telco, and browsers to service providers.

While the Internet architecture imaginary is still alive, it can be doubted whether the it will produce anything else than interoperation between technologies of large transnational corporations

Abbate, Janet. 1999. *Inventing the Internet*. Inside Technology. Cambridge, MA etc: The MIT Press.
Aboba, Bernard, Elwyn B. Davies, and Elwyn Davies. 2007. “RFC4924 - Reflections on Internet Transparency.” https://tools.ietf.org/html/rfc4924.
Braman, Sandra. 2009. “Internet RFCs as Social Policy: Network Design from a Regulatory Perspective.” *Proceedings of the American Society for Information Science and Technology* 46 (1): 1–29. https://doi.org/10.1002/meet.2009.1450460254.
———. 2013. “Laying the Path: Governance in Early Internet Design.” *Info* 15 (6): 63–83. https://doi.org/10.1108/info-07-2013-0043.
———. 2016. “Instability and Internet Design.” *Internet Policy Review* 5 (3).
Bush, Randy, and David Meyer. 2002. “RFC3439 - Some Internet Architectural Guidelines and Philosophy.” https://tools.ietf.org/html/rfc3439.
Carey, James. 1989. *Communication as Culture: Essays on Media and Culture*. New York, NY: Unwin Hyman.
Carpenter, Brian E. 1996. “RFC1958 - Architectural Principles of the Internet.” https://tools.ietf.org/html/rfc1958.
———. 2000. “RFC2775 - Internet Transparency.” https://tools.ietf.org/html/rfc2775.
Carpenter, Brian E., and Scott W. Brim. 2002. “RFC3234 - Middleboxes: Taxonomy and Issues.” https://tools.ietf.org/html/rfc3234.
Carr, Madeline. 2013. “Internet Freedom, Human Rights and Power.” *Australian Journal of International Affairs* 67 (5): 621–37. https://doi.org/10.1080/10357718.2013.817525.
Castells, Manuel. 2009. *Communication Power*. Oxford, UK; New York: Oxford University Press. http://public.eblib.com/choice/publicfullrecord.aspx?p=472226.
———. 2011. “A Network Theory of Power.” *International Journal of Communication* 5 (0): 15.
Cath, Corinne, and Luciano Floridi. 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. https://doi.org/10.1007/s11948-016-9793-y.
Cowhey, Peter F., Jonathan D. Aronson, and John Richards. 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, Jon. 2007. “Net Neutrality: The Technical Side of the Debate: A White Paper.” *SIGCOMM Comput. Commun. Rev.* 37 (1): 49–56. https://doi.org/10.1145/1198255.1198263.
Den Hertog, J. A. 2010. “Review of Economic Theories of Regulation.” *Discussion Paper Series/Tjalling C. Koopmans Research Institute* 10 (18).
DeNardis, Laura. 2009. *Protocol Politics: The Globalization of Internet Governance.* Cambridge, MA: MIT Press.
———. 2014. *The Global War for Internet Governance*. New Haven: Yale University Press.
Dolata, Ulrich, and Jan-Felix Schrape. 2018. *Collectivity and Power on the Internet: A Sociological Perspective*. SpringerBriefs in Sociology. Springer International Publishing. //www.springer.com/de/book/9783319784137.
Flichy, Patrice. 2007. *The internet imaginaire*. Cambridge, Mass.: MIT Press.
Freed, Ned. 2000. “Behavior of and Requirements for Internet Firewalls.” https://tools.ietf.org/html/rfc2979.
Galloway, Alexander R. 2006. *Protocol: How Control Exists After Decentralization*. New Ed edition. Cambridge, Mass.: MIT Press.
Gibson, James J. 1978. “The Ecological Approach to the Visual Perception of Pictures.” *Leonardo* 11 (3): 227–235.

Guest, Greg, Kathleen MacQueen, and Emily Namey. 2012. *Applied Thematic Analysis*. 2455 Teller Road, Thousand Oaks California 91320 United States: SAGE Publications, Inc. https://doi.org/10.4135/9781483384436.

Hanseth, Ole, and Eric Monteiro. 1997. “Inscribing Behaviour in Information Infrastructure Standards.” *Accounting, Management and Information Technologies* 7 (4): 183–211. https://doi.org/10.1016/S0959-8022(97)00008-8.

Hayes, David A., Jason But, and Grenville Armitage. 2008. “Issues with Network Address Translation for SCTP.” *ACM SIGCOMM Computer Communication Review* 39 (1): 23. https://doi.org/10.1145/1496091.1496095.

Huston, Geoff. 2017. “The Internet’s Gilded Age.” *The ISP Column* (blog). March 2017. http://www.potaroo.net/ispcol/2017-03/gilding.html.

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.” https://tools.ietf.org/html/rfc3724.

Internet Society. 2012. “Internet Invariants: What Really Matters.” *Internet Society* (blog). February 3, 2012. https://www.internetsociety.org/internet-invariants-what-really-matters/.

———. 2013. “The Value of Openness for a Sustainable Internet.” *Internet Society* (blog). 2013. https://www.internetsociety.org/resources/doc/2013/the-value-of-openness-for-a-sustainable-internet/.

Isin, Engin, and Evelyn Ruppert. 2015. *Being Digital Citizens*. London, UK ; New York: Rowman & Littlefield International.

Jasanoff, Sheila. 2004. “‘Ordering Knowledge , Ordering Society.” In *States of Knowledge: The Co-Production of Science and the Social Order.*, edited by Sheila Jasanoff. Vol. International library of sociology. London: Routledge.

Jasanoff, Sheila, and Sang-Hyun Kim. 2015. *Dreamscapes of Modernity: Sociotechnical Imaginaries and the Fabrication of Power*. University of Chicago Press.

Johnson, David R., Susan P. Crawford, and John G. Jr Palfrey. 2004. “The Accountable Internet: Peer Production of Internet Governance.” *Virginia Journal of Law & Technology* 9: 1.

Kahin, Brian, James Keller, and Harvard Information Infrastructure Project. 1995. *Public Access to the Internet / Ed. by Brian Kahin and James Keller*. Cambridge, Mass [etc.: The MIT Press.

Kelty, Christopher. 2005. “Geeks, Social Imaginaries, and Recursive Publics.” *Cultural Anthropology* 20 (2): 185–214.

Lessig, Lawrence. 2006. *Code 2.0*. New York: Basic Books. http://codev2.cc/download+remix/Lessig-Codev2.pdf.

Litan, Robert E., and Alice M. Rivlin. 2001. “Projecting the Economic Impact of the Internet.” *American Economic Review* 91 (2): 313–317.

Losey, James, and Sascha D. Meinrath. 2016. “In Defense of The Digital Craftsperson.” *Journal of Peer Production*, no. 9.

Mansell, Robin, and Michele Javary. 2002. “Emerging Internet Oligopolies: A Political Economy Analysis.”

Mathew, Ashwin J. 2014. *Where in the World Is the Internet? Locating Political Power in Internet Infrastructure*. Berkeley: University of California. https://www.ischool.berkeley.edu/research/publications/2014/where-world-internet-locating-political-power-internet-infrastructure.

McLuhan, Marshall. 2013. *Understanding Media: The Extensions of Man*. Critical edition. Gingko Press.

Morris, John, and Alan Davidson. 2003. “Policy Impact Assessments: Considering the Public Interest in Internet Standards Development.” SSRN Scholarly Paper ID 2060656. Rochester, NY: Social Science Research Network. https://papers.ssrn.com/abstract=2060656.

Mosco, Vincent. 2005. *The Digital Sublime: Myth, Power, and Cyberspace*. The MIT Press.

Musiani, Francesca, Derrick L Cogburn, Laura DeNardis, and Nanette S Levinson. 2015. *The Turn to Infrastructure in Internet Governance*.

Poullet, Yves. 2001. “How to Regulate Internet: New Paradigms for Internet Governance Self-Regulation: Value and Limits.” *Variations Sur Le Droit de La Soci Ét é de l’information*, 79–114.

Price, Monroe, and Stefaan Verhulst. 2000. “The Concept of Self-Regulation and the Internet.” *Protecting Our Children on the Internet: Towards a New Culture of Responsibility*, January, 133–98.

Rachovitsa, Adamantia. 2017. “Rethinking Privacy Online and Human Rights: The Internet’s Standardisation Bodies as the Guardians of Privacy Online in the Face of Mass Surveillance.” SSRN Scholarly Paper ID 2911978. Rochester, NY: Social Science Research Network. https://papers.ssrn.com/abstract=2911978.

Rogers, Michael, and Grace Eden. 2017. “Digital Citizenship and Surveillance| The Snowden Disclosures, Technical Standards, and the Making of Surveillance Infrastructures.” *International Journal of Communication* 11 (0): 22.

Russell, Andrew L. 2014. “Open Standards and the Digital Age: History, Ideology, and Networks.” *Afterimage* 42 (3): 39.

Rüth, Jan, Ingmar Poese, Christoph Dietzel, and Oliver Hohlfeld. 2018. “A First Look at QUIC in the Wild.” *ArXiv:1801.05168 [Cs]*, January. http://arxiv.org/abs/1801.05168.

Saltzer, J., D. Reed, and D. Clark. 1984. “End-to-End Arguments in System Design.” *ACM Transactions on Computer Systems (TOCS)* 2 (4): 277–288. https://doi.org/10.1145/357401.357402.

Saunders, Benjamin, Jenny Kitzinger, and Celia Kitzinger. 2015. “Anonymising Interview Data: Challenges and Compromise in Practice.” *Qualitative Research* 15 (5): 616–32. https://doi.org/10.1177/1468794114550439.

Sennett, Richard. 2008. *The Craftsman*. Yale University Press.

Shapiro, Carl, and Hal R. Varian. 1998. *Information Rules: A Strategic Guide to the Network Economy*. Boston, Mass: Harvard Business Review Press.

Singh, J. P. 2008. *Negotiation and the Global Information Economy*. Cambridge University Press Cambridge.

Srisuresh, Pyda, and Matt Holdrege. 1999. “RFC2663 - IP Network Address Translator (NAT) Terminology and Considerations.” https://tools.ietf.org/html/rfc2663.

Star, Susan, and Karen Ruhleder. 1994. “Steps towards an Ecology of Infrastructure: Complex Problems in Design and Access for Large-Scale Collaborative Systems.” In , 253–264. CSCW ’94. ACM. https://doi.org/10.1145/192844.193021.

Stewart, Randall. 2007. “RFC4960 - Stream Control Transmission Protocol.” https://tools.ietf.org/html/rfc4960.

Taylor, Tom, Hanns Juergen Schwarzbauer, Malleswar Kalla, Lixia Zhang, Ken Morneault, Ian Rytina, Randall R. Stewart, Chip Sharp, and Qiaobing Xie. 2000. “RFC2960 - Stream Control Transmission Protocol.” https://tools.ietf.org/html/rfc2960.

Ten Oever, Niels, and Kathleen Moriarty, eds. 2018. “The Tao of IETF: A Novice’s Guide to the Internet Engineering Task Force.” *IETF*. https://www6.ietf.org/tao.

Thomas, Graham, and Sally Wyatt. 1999. “Shaping Cyberspace—Interpreting and Transforming the Internet.” *Research Policy* 28 (7): 681–698.
Tuexen, Michael, Irene Ruengeler, and Randall R. Stewart. 2013. “RFC7035 - SACK-IMMEDIATELY Extension for the Stream Control Transmission Protocol.” https://tools.ietf.org/html/rfc7053.
Van Schewick, Barbara. 2012. *Internet Architecture and Innovation*. Mit Press.
Wyatt, Sally. 2004. “Danger! Metaphors at Work in Economics, Geophysiology, and the Internet.” *Science, Technology, & Human Values* 29 (2): 242–61. https://doi.org/10.1177/0162243903261947.
Zittrain, Jonathan. 2008. *The Future of the Internet–and How to Stop It*. Yale University Press.