Whither the public Internet?

William Lehr[1]
David Clark
Steve Bauer
Arthur Berger
Philipp Richter

Table of Contents

Abstract

This paper addresses the question of what should be meant by the phrase *the Internet*, or equivalently, *the public Internet*. Since its origins in the 1960s, the Internet has changed significantly in terms of the networks and technologies, services that are supported, and industry players who participate in the Internet ecosystem, growing in global economic and social importance. This paper discusses the changing role of the network operators and service providers that support the public Internet, and the relationship between what might be considered the public Internet and other elements of the global infrastructure. Herein, we explain why the quest for a single definition for *the Internet* is likely a fool's errand, while recognizing that there remains a need to understand what the concept means in light of its growing importance as a topic of policy concern. Instead of a definition, we propose a three-lens framework through which to evaluate technical

---

[1] William Lehr <wlehr@mit.edu>, corresponding author; David Clark <ddc@csail.mit.edu>; Steve Bauer <bauer@mit.edu>; Arthur Berger <awberger@mit.edu>; and Philipp Richter <richterp@csail.mit.edu>.

> and market trends that are changing the Internet and assist in determining whether those changes are or are not consistent with what we view as worth preserving in that which we refer to as *the Internet*.

1. Introduction

The Internet is touted by many as essential infrastructure.[2] That assertion, however, begs the question of exactly what the Internet is. What do we mean or should we mean when we say "the Internet," (or more specifically, "the public Internet" -- terms that for our purposes here may be used interchangeably)? One strategy for answering this question was to look for prior definitions in the law, on the Web, or in the academic literature. What that exercise teaches one is that there is no single definition that everyone agrees on beyond something banal like "a network of networks," and suggests that the exercise of finding any single best definition may be a fool's errand. On the other hand, in an age when there is an active debate over the need to regulate the Internet, and numerous nations are identifying the Internet as essential infrastructure and asking whether Internet access should be a basic right for all citizens, it seems important to consider what one might mean by the term, "the Internet." Moreover, for this to be an interesting question, it is more important to ask where we think the Internet may be headed and whether we are happy with the current trajectory.

In this paper we identify technical and market trends that are changing the Internet and highlight why a simple definition is neither feasible nor desirable. Instead, we propose a series of lenses with which to evaluate how the Internet is changing to assist in determining whether those changes are or are not consistent with what we view as worth preserving in that which we refer to as "the Internet."

Much of the groundwork for understanding why the term "the Internet" (or, "the public Internet") is difficult to define uniquely was covered in an excellent paper by Cannon (2003). Cannon focused his analysis on the challenge of finding an appropriate definition in U.S. statutes and regulatory codes and noted the lack of consistent definitions. As Cannon explains, having definitions in the law is important to reduce uncertainty and enhance trust that the law will be enforced as expected. Multiple definitions in overlapping statutes cause confusion. Cannon's final conclusion, with which we agree, is that the law might be better if it could avoid needing to define the Internet at all, focusing on a class of network that should receive consistent treatment rather than on a specific network called "the Internet"; however, Cannon also concludes that avoiding an attempt to define the term, "the Internet" is not completely acceptable because the Internet is distinct with respect to existing U.S. statutory codes and legal frameworks such as the First Amendment and the Communications Act.[3]

---

[2] As of 2016, the UN has identified Internet access as a basic human right (see Howell and West, 2016).

[3] Cannon points out that in assessing First Amendment free speech claims, the Supreme Court has noted that the nature of the medium should be taken into consideration and points to special features of the Internet that distinguish it from other expression media. Cannon also notes the relevance of determining whether a service such as broadband access to the Internet is a telecommunications service in light of Title II of the U.S. Communications Act. At the time, the FCC had classified broadband access as an "information service," thereby exempting it from Title II regulation, but the debate was far from over and continues to this day.

Our concern here is broader than Cannon's since we are not solely concerned with the challenges U.S. lawyers confront when trying to craft or interpret legal codes based on conflicting or imprecise definitions. Even if there were no lawyers, stakeholders may need to have a set of criteria to decide whether a particular network, trend, or behavior was consistent with their concept of what the Internet is, where it is going, or where it should go. Having a shared understanding of the concept ought to help in focusing policy discussions.

In the following we present several possible criteria that might characterize the Internet, in the hope that those engaged in discussing where they think the Internet is going or should go will find it useful. We organize our criteria through three lenses, reflecting the perspectives of three sets of actors that we think are most relevant: (1) abstract architectural; (2) network complementors; and (3) customer experience. Although we believe this tri-part distillation is useful, our focus in presenting it is to use it to elucidate how the Internet has been changing and some of the policy concerns this may raise. We also wish to reaffirm and expand on Cannon's earlier conclusion that seeking to define or locate policy too directly on any single or overly precise definition of the Internet is undesirable in light of the need to allow the Internet to continue to evolve, potentially in ways we cannot fully anticipate today. Moreover, while we think it is reasonable to expect to have conceptual frameworks or filters that would provide a basis for distinguishing many behaviors, actors, services, or events as either "Internet" *or* "not-Internet," there will also be gray-area cases that will need to be resolved on a context-specific basis. For example, determining whether a particular service, device, or traffic is "on the public Internet" or not may depend on the context of the question, as we will elaborate further below.

In titling this paper, we include "whither" to signal our interest in considering how the Internet has changed and the trends that are driving those changes, as well as some of our normative views on what we believe is essential and worth preserving as the Internet continues to evolve. We include "public" to signal that we recognize that not all IP networks may even be connected to the Internet and even when connected, may more appropriately be regarded as private;[4] although absent a clear legal demarcation point, identifying where the boundary ought to be (by physical location, by who owns or makes control decisions, by layer, by the source or destination of the traffic, or by some other method) is open to debate as we have discovered in our collective efforts to write this paper. We also include "public" to signal that there is a public interest in the Internet. Some portions of it may be publicly owned/operated (or some might argue, should be) and policymakers have identified Internet access as essential infrastructure which governments have an obligation to ensure all citizens have reasonable access to. Absent a public interest, the Internet would not be worthy of on-going policy concern.

The balance of this paper is organized into four sections. In Section 2, we review some of the conflicting definitions of the Internet available in the literature and highlight why these do not

---

[4] For example, a standalone factory IP network that is not connected to the Internet is not part of any "network-of-networks." Moreover, as we will explain further below, we reject defining the Internet even in terms of the specific protocols that have been most closely associated with it through its history to date. Although, while it is clear that a standalone factory IP network (which may more appropriately be referred to as an "intranet") is not part of the Internet, it is part of the larger Internet ecosystem since it contributes to the demand and usage cases for the hardware and software solutions on which the Internet depends.

adequately address our concerns. In Section 3, we present our three-lens framework. Section 4 uses these three lenses to discuss important trends that are changing the Internet and Section 5 offers concluding thoughts.

2. Defining "the Internet" – a fool's errand?

If one searches the Web for definitions of "the public Internet," one of the prominent results that comes up is from PC Magazine's Encyclopedia, which defines it as "a publicly accessible system of networks that connects computers around the world via the TCP/IP protocol."[5] It goes on to note that a "public Internet implies that everyone has access," and that the "term mostly refers to the global Internet" which "comprises a billion Web, email and related servers in more than 100 countries."[6] The U.S. code of statutes variously defines the term "Internet" to mean "the international computer network of both Federal and non-Federal interoperable packet switched data networks,"[7] and "the combination of computer facilities and electromagnetic transmission media, and related equipment and software, comprising the interconnected worldwide network of computer networks that employ the Transmission Control Protocol/Internet Protocol or any successor protocol to transmit information."[8]

As Cannon (2003) noted and we agree, defining the Internet in terms of a specific subset of protocols is problematic because not all of the networks that are part of the Internet implement the same set of protocols in the same way and the suite of Internet-related protocols has continuously evolved. The specific protocol choices may be irrelevant to the nature of the end-user experience and relying on a definition anchored in a particular specification is contrary to policymakers' aspirations for technically neutral regulations.

Another common approach for interpreting what is meant by "the Internet" is in historical terms. Viewed from this perspective, the Internet is the global network of interconnected endpoints, web of data, and services that evolved from the U.S. government-sponsored ARPAnet network that was originally developed in the 1960s to support communications between mainframe computers and end-user terminals managed by diverse government and university computing facilities.

While the essential forwarding protocol of the Internet (the Internet protocol itself, or "IP") has remained stable since the early 1980's, much has changed with respect to reach and services. The Internet was commercialized and grew to mass market prominence with the growth of personal computing, the rise of email and local area networking, and the invention of the World Wide Web in the 1990s. After 2000, the rise of broadband Internet access greatly expanded the range of applications and interactive multimedia services that could be supported over the Internet. With

---

[5] See https://www.thefreedictionary.com/Public+Internet (visited May 27, 2018).

[6] See https://www.pcmag.com/encyclopedia/term/59572/public-internet (visited May 27, 2018). Note that the definition uses the term "mostly" and focuses on "public" which differentiates the "public Internet" from private or "intranet" networks (where access is restricted to a closed user group).

[7] See 47 U.S.C. §230(f)(1).

[8] See 47 U.S.C. §231(e)(3). For further statutory definitions of "Internet" and a complementary discussion of the multiple definitions that have been advanced, see http://www.cybertelecom.org/notes/internet_definition.htm (visited May 27, 2018).

advances in mobile and wireless networks and the rise of smartphones and other post-PC connected devices like eReaders, tablets, and Internet appliances (e.g., thermostats, health trackers, etc.), options for connecting to the Internet expanded significantly. Concurrently, the growth of Internet platforms for eCommerce like eBay, Amazon, and AirBnB; for social networking like Facebook, Twitter, and YouTube; and on-line chat, bloggers, and other on-line content creation and sharing platforms have expanded the range of content and services accessible via the global Internet.

From this historical perspective, the Internet is this global conglomeration of networked services, content, devices, and network end-points that has emerged and continues to evolve and morph. Such an approach captures much of what a reader of a general interest business magazine might understand as "the Internet." But like the attempt to anchor our understanding in terms of the underlying technology, this approach (while demonstrably useful),[9] does not allow one to discern how far one might move from what the Internet was in the past to where it is going before it would cease to be the Internet and become something else.

In the next section, we propose our three lenses to facilitate structured discrimination between trends, services, behaviors, or network futures that are either consistent with "the Internet" *or* not. As noted, we believe having this capability is important as we engage in a global debate over the future of "Internet policy," or (telecommunications or digital economy policy) since however these debates proceed (whether in the court room or in forums of public opinion) the term "the Internet" will be commonly encountered. Thus, we believe having a framework to understand what we might mean by the concept will be important.

In presenting our lenses, we do not seek to dismiss the other "definition" approaches noted earlier nor claim that our proposed framework is either unique or the best. Our lenses arise from a range of concepts and ideas we have been collectively grappling with within our research group at MIT and with colleagues in the Internet research and policy communities for a number of years.[10] Our goal is to spark further discussion and consideration of what our collective aspirations and concerns should be for "the Internet," a discussion that we do not expect to end anytime soon.

3. The Internet through three lenses

Hopefully we have convinced readers that seeking a single definition of the Internet is complicated and may be unnecessary, while still keeping readers interested in developing a refined conceptual framework for determining whether particular services, actors, or behaviors are consistent with the concept of "the Internet" *or* not.

---

[9] When used in the trade press or normal conversation between non-experts and experts alike, virtually no one other than someone asking the questions we ask here would scratch her head and say "what does she mean by 'the Internet'?". This demonstrates the practical use of loose conceptualizations.

[10] For example, see Claffy & Clark (2014, 2015), Lehr & Sicker (2017, 2018), Lehr, Clark and Bauer (2013), and Faratin, Clark, Bauer, Lehr, Gilmore, and Berger (2008) for various takes on how to think about the Internet in relation to communications policy.

Here we explain our lenses for characterizing the Internet (weaker than "defining") to provide such a framework. Our three lenses are: (1) Abstract architectural; (2) Network complementors; and (3) Customer experience.

3.1. Abstract Architectural

When asked to characterize the Internet, a network engineer might answer "a set of networks linked by a global address space and a common forwarding protocol that allow connected end-nodes to communicate by the direct exchange of packets."

The individual networks that make up the Internet are referred to as Autonomous Systems (ASes). An Autonomous System can be defined as a region of the Internet associated with one or more blocks of allocated addresses (called "routing prefixes") under the control of one or more network operators on behalf of a single administrative entity or domain that presents a common, clearly defined routing policy to the Internet.[11] The Autonomous Systems exchange routing information so that the addresses allocated to each can be reached by all the other ASs, and they interconnect so that traffic can flow according to that routing information.

For a network architect, the essential features of the Internet include that (1) it was designed as an *end-to-end (e2e)* peer-to-peer network, which by design was intended to be agnostic about the information content of the packets; (2) it was based on a *layered architecture*; and, (3) it provides a *minimalist spanning layer*, with limited in-network state or intelligence, that allowed it to be implemented on a wide array of physical network infrastructures, enabling end-to-end interoperability and interconnections across diverse types of networks.

By "layered architecture," we want to highlight the generality and scope of the basic architecture. The Internet as a packet transport system exists in conjunction with other layers that it interacts with. Early pictures of the Internet's architecture characterized it as the narrow waist of an hourglass.[12] The term "spanning layer" captures the idea that the uniform packet forwarding service "spans" all the lower layer variants of network technology to provide the end-to-end packet transport. The protocols that implemented the early Internet provided a spanning layer that allowed diverse higher-level functions in the upper part of the hourglass, including applications, to communicate across the networks implemented by diverse physical network infrastructures in the lower part of the hourglass. The upper applications and the lower infrastructures were insulated from needing to know about each other by this layer.

An economist considering the Internet abstractly might describe it as a general-purpose platform for electronic communications, eCommerce, and digital services. Such a platform may operate as a *General Purpose Technology* (GPT), which is an input to the production of many different kinds of goods and derives its economic value from the values of the goods it contributes to producing.[13]

---

[11] See https://en.wikipedia.org/wiki/Autonomous_system_(Internet) for a related definition.

[12] For example, see Figure 2.1, page 53, in CSTB (1994).

[13] See Bresnahan and Trajtenberg (1995). The identify steam engines and electric motors as earlier examples of GPTs, and computer processors and semiconductors as GPTs today. What distinguishes a GPT is its pervasive use throughout the economy and its ability to generate productivity improvements that cascade

Moreover, innovations in a GPT create spillovers in inducing productivity innovations in downstream markets. In considering the role of the Internet as a GPT, its key features include that its ownership and control is distributed and decentralized; and that it is an open platform with low entry costs for collaborating in the development and deployment of new services and applications that has allowed the Internet to support significant edge-based, permission-less innovation. Zittrain (2006) characterized this feature as the "Generative Internet."[14]

The fact that ownership and control is distributed and decentralized is important because it means that the character of the Internet is not under the control of a single powerful actor. The Internet serves as a platform for the application developers that sit at the layer above the forwarding function. It is important for the health of that developer ecosystem that this platform layer be stable[15] and open. Having different parts of the Internet built and operated by different actors imposes competitive discipline on the providers, and the necessity to coordinate the evolution of the platform among all the providers is a force for stability.

Emphasizing the ability of networks such as the Internet to reduce the marginal cost of reproducing and distributing digital information to (near) zero, a number of analysts have characterized the Internet as a "digital commons," focusing on its shared ownership that renders it amenable to analysis as a Common Pool Resource (CPR).[16]

From this perspective, defining what constituted "the Internet," or equivalently, "the public Internet" would focus on the appropriate definition of the "public" and would be aspirational in the sense that it would focus on defining the desirable features that such a public Internet should provide.[17]

In keeping with the notion of open access, the public Internet would be defined broadly, presumably to include any end-user with the appropriate technology (hardware and software) to allow the end-user to establish an IP connection. Indeed, proponents of the Internet as a platform for the "global digital commons" envisioned it as enabling "an unprecedented number of people

---

through the economy, resulting in increasing returns to scale. Economists have also characterized the Internet as a GPT (e.g., see Harris (1998) or Clark, Qiang, and Xu (2015)).

[14] Zittrain described the Internet as a "network that no one in particular owns and that anyone can join."

[15] By stable, we do not mean never-changing. The Internet needs to be able to evolve, but too rapid or uncoordinated changes to the basic platform would render it inadequately stable. The stability of the Internet as a platform, however, does pose a significant challenge for enabling desirable evolution. Yoo (2016) offers an interesting perspective on how to the architecture of the Internet should be interpreted to be consistent with the need to allow the Internet to evolve.

[16] This follows in the shoes of Elinor Ostrom's (1990) Nobel-winning work on the governance of CPRs. When viewed as a digital commons, the focus is on the Internet's role as a platform for producing, providing access to, and sharing digital content. See Hess (2000), Hess & Ostrom (2007), Hofmokl (2010), De Rosnay & Le Crosnier (2012), and Sowell, J. (2015).

[17] There are many candidates for Internet aspirations that could be articulated. See Claffy and Clark (2015) for one list.

around the world to freely express themselves and come into contact with the ideas and opinions of others."[18] This is a vision of the public Internet as a forum to support global democracy.

We refer to both these engineering and economic characterizations (which are intentionally presented as less than definitions) of the Internet as "abstract architectural" (for want of a better term) to signal the sense in which we intend them to be distillations at an abstract and architectural level of what we as co-authors with backgrounds in engineering and economics consider are the essential features of the Internet that are worth preserving.

Refining the earlier discussion further (Table 1), the engineering features we think are most important and relevant are: that the Internet was designed as a (1) layered network architecture that supports (2) end-to-end connectivity at the packet layer, exploiting a (3) global address space to support packet forwarding across multiple (4) interconnected ASes with (5) global reach via (6) a decentralized inter-AS routing protocol and (7) adhering to a shared set of standardized protocols for forwarding and inter-AS routing.

Notice that we drop any mention of IP/TCP since those protocols may be replaced. By "packet level connectivity" we mean that the network should be capable of supporting packet forwarding across links and networks with heterogeneous quality of service (QoS) capabilities; the term "packet" means collections of digital bits and are signaling that the fundamental communications paradigm strives to be agnostic regarding the semantics of what the bits are and how QoS is assured.

By "global address space" we mean only that there are addresses that are capable of being uniquely resolved within the "network of networks," while recognizing that the nature of the addressing and identification scheme may change (e.g., from IPv4 to IPv6). By focusing on "interconnected ASes," we signal that the Internet is a "network of networks" which are not under centralized control by any single management entity.

By "global reach" we signal that the Internet should be scalable both in terms of the number of end-nodes that could be included and that it should not be limited in geo-scope. That is, like today's Internet, it should be able to span sovereign boundaries around the globe (and maybe even extend into space).[19] The routing protocol is decentralized in the sense that ASes individually decide what reachability information is shared with each neighboring AS, and whereby the self-interest of individual ASes results in global reachability between end point pairs on the public Internet.

---

[18] See Lambino (2010) who laments the rise of Facebook and other service providers that are seeking to use the Internet to commercialize consumer data, thereby posing a threat to " the realization of a global digital commons, one envisioned to enable an unprecedented number of people around the world to freely express themselves and come into contact with the ideas and opinions of others."

[19] As distances become large, the speed of light imposes constraints on attainable latency goals. An advantage of the Internet's legacy architecture is its inherent robustness. Enhancements to the Internet have included expanding support for delay or disruption tolerant networking, the need for which becomes immediately obvious at planetary-scale transmission distances, but also arises commonly with wireless and in other typical communication contexts in which we may wish to access the Internet. There is a large body of work on delay tolerant networking (e.g., see Fall, 2003).

Finally, by "adhering to a shared set of standardized protocols" we want to indicate the mechanism by which the networks that wish to participate in its network of networks manage the technical aspects of their interconnection and interoperability. While not every AS needs to implement every protocol in the shared set of standardized protocols, it is by subsets of those protocols that the connectivity fabric of the Internet is sustained. Those protocols are shared property of the Internet community, and are managed today principally by the Internet Engineering Task Force (IETF).[20] However, the governance structure may change (e.g., the IETF could cease to exist), and the Internet may still remain since there are many ways in which the networking community that sustains the shared set of standards might organize itself.

The economic features we think are most important and relevant include that the Internet should be: (1) a general purpose platform for electronic communications, that serves as (2) a platform for markets and supports (3) open access and (4) permission-less innovation, (5) implemented across networks for which ownership and management is decentralized and distributed.

By general purpose platform, we wish to signal that the Internet is basic economic infrastructure that is broadly used across society and the economy, rather than for a specialized or niche set of services or functions. In economic terms, we regard the Internet as a GPT.

By open access, we mean that there are no a priori restrictions that limit connectivity to a closed user group, or on what or who may access the Internet and imply that the entry barriers to access are generally low. We also mean to capture the economic sense in which access to the Internet is not under the control of any single economic entity.[21] The distinction between an *Intranet* vs. *Extranet* vs. (the) *Internet* captures part of this notion: intranets connect users within a private network (e.g., a single enterprise); extranets connect users within a restricted community (e.g., suppliers and customers); while the Internet is accessible to the general public. As noted earlier, deciding where an Intranet ends and the Internet begins is not always easy in practice or concept. Extranets represent a gray area between private Intranets and the public Internet.

By permission-less innovation, we mean that users are able to launch new applications, sub-communities, and services virally. Experimentation with new technologies and business models should be allowed by default.

The fact that control and ownership is distributed and decentralized signals that there is no single entity that has decision-making authority over the Internet's design or operation.[22] Additionally,

---

[20] See http://ietf.org/about/.

[21] In this, we are echoing Bresnahan (2015) when he says "openness is an economic organization concept. It refers to the availability and control of information about interface standards and to the role of a platform sponsor as a gatekeeper."

[22] Nominally, the IETF is the institution that the networking community looks to as the principal entity in charge of coordinating the design of the protocols that manage the Internet, but its authority is far from absolute and depends on the consensus cooperation of the members. Of course, the same could be said of any international institution, but it is worth noting that the IETF does not even have the status of an international institution secured by a multilateral treaty among sovereign nations, like the ITU.

the fact that ownership is distributed and decentralized signals that the residual claimants for the economic value created by the Internet and the investment and operating costs incurred to create the Internet are decentralized across markets, enterprises, and industries. That implies that one should not assume that all of the agents' incentives and interests are well-aligned and that strategic behavior is likely to be a factor in how the Internet evolves. The future of the Internet depends on a market-place tussle.[23]

The current Internet is not the only possible approach to meeting the criteria in Table 1. We refer in this paper to the Internet (with a capital 'I') to distinguish it from other possible designs for an internet (with a lower case 'i') that could have emerged in the past or might emerge in the future. As we will discuss in section 3.2, there are other global networks in operation today, often using the same protocols (most critically, the Internet protocol). There have always been alternative networks built using the Internet protocols, such as the intranets mentioned above. Specialized networks exist today to provide interconnection services between enterprise and cloud, and meet other specific needs for packet carriage. None of them meet all the criteria in Table 1, so none of them would qualify, by our definition, as examples of an internet. But it is important to remember that there could be many specific examples of an internet that would fit our abstract architectural criteria: the Internet is just the example that today represents the class.[24]

Table 1: Abstract Architectural Criteria for Characterizing the Internet

Network Engineering	Economic
(1) layered architecture	(1) General Purpose Platform
(2) end-to-end packet connectivity	(2) Markets
(3) global address space	(3) Open Access
(4) interconnecting multiple ASes	(4) Permission-less Innovation
(5) global reach	(5) Decentralized, distributed ownership & control
(6) inter-AS routing protocol	
(7) shared set of standardized protocols	

3.2. Network Complementors

The Internet derives much of its value because of the diverse applications and uses it supports, and because of its role as a platform for permission-less innovation. Much of the focus of policy debates has been on the need to ensure that the Internet remains an open platform for application and content providers, or so-called "edge providers."

Since these actors exploit the Internet as a platform for innovation, we refer to them as platform complementors. Viewed from this perspective, the Internet offers a set of

---

[23] See Clark et al. (2002).

[24] https://en.wikipedia.org/wiki/Capitalization_of_%22Internet%22 contains an extensive discussion of 'I' vs. 'i'. However, that entry does not contemplate the need to identify other possible constructions that match our abstract architectural criteria.

functionalities/capabilities (at the wholesale level[25]) to other participants in the larger business ecosystem that use the Internet as a platform for the deployment, implementation, and operation of their applications and services. In the early days of the Internet, the application development ecosystem was very simple. Application code ran on end-nodes that were connected by the forwarding capability of the Internet. Today the application development ecosystem includes many more platform assets than just basic Internet transport, including cloud services (dynamic access to computing and storage), Content Delivery Networks (CDNs), and other ancillary services such as fault recovery, cybersecurity monitoring, access management, etc.) The application developer also has access to packet transport services that are provided by networks that may be based on the same technology as the Internet (the Internet protocol), but which are not a part of the Internet as we defined it through the abstract architecture lens (example include cloud-internal networks and direct interconnections with such cloud networks, which we discuss in Section 4.3). These other networks, many of which have global reach, do not share a common global address space, and do not interconnect with the Internet to provide general packet level forwarding.

Different application developers may exploit these various platform elements in different ways. For example, Netflix started out by using the Internet as a way to connect with customers, using a Web-based interface to enable end-users to connect with Netflix's offerings. The Web-based interface allowed Netflix to present end-users with a directory to its large library of content and for end-users to select which content they would like to rent. To deliver the content, Netflix chose a very high-latency, but also very high data-rate transmission network – the U.S. postal service for the delivery of DVDs from Netflix's warehouses to customers in Netflix's custom mailer envelopes. In recent years, Netflix has shifted to delivering its media content via its customers' broadband access services. In the earlier delivery model, Netflix bundled the delivery costs into its service price – customers were shipped DVDs with pre-paid return envelopes. In the current model, Netflix relies on customers to provide the broadband access connection that allows those customers to connect with Netflix's content servers. When Netflix began to offer its streaming content services via the Internet it contracted with third-party CDN providers to manage the delivery of the traffic from Netflix's servers to Netflix customers. Today, Netflix traffic is so high volume that it makes more sense for Netflix to manage its own CDN, which is something Netflix does, along with a growing number of other very large content providers such as Google (YouTube) and Facebook. When providers build their own CDN, they are depending on yet another element in the platform ecosystem: colocation facilities where providers can rent space for server equipment and interconnect to the Internet.

The management of these CDNs provides an illustration of the options open to application developers today for data transport. CDN servers must connect to the public Internet so that they can deliver data to customers on demand. As well, the servers need to have a way to communicate with the original source of the content they hold, so that this content can be loaded into the server. In earlier times, the only way this connection could be made was also over the public Internet. Today, the designer of the CDN has the option of either using the public Internet for this connection, or one of the other global networks based on IP mentioned above. That CDN designer could also build his own global network to update and manage his CDN servers. Google has built

---

[25] We use "wholesale" here to refer to business customers who make use of the Internet as an intermediate input to their offering of a "retail" service of final good that they provide to their end-users or consumers.

such a network—an IP-based network with global reach only used for traffic internal to Google, which connects the various Google data centers, and could be used to connect to some of its CDN servers that host YouTube content. These other global networks may traverse physical infrastructure (fiber cables, network servers, routers) that carry both general Internet traffic and those flows; in other cases, the traffic may travel via dedicated physical paths that are separate both logically and physically from the Internet.

This example illustrates a number of features of the relationship between platform complementors and the Internet. First, the complementors view the Internet as only one component of the available platform assets, which for some uses may have non-Internet alternatives. The complementor can be expected to pick and choose whether to use the Internet or an alternative when the alternative is appropriate, and in many cases, the complementors may choose to use (or implement) non-Internet assets because those are more capable, less expensive, more secure, or provide the complementors with more control of their business and its relationship with their customers. The scope of the platform assets available today to platform complementors (innovators) is much greater than just the public Internet we characterized through the lens of abstract architecture, and for a regulator concerned with the aspiration of third-party innovation, a too-narrow focus on only the public Internet itself may miss issues that might in the future impair the goal of innovation.

There are many protocols used as part of the Internet. The most central is the Internet protocol itself (IP), which provides the addressing information necessary to the forwarding function.[26] TCP is used end-to-end when reliable delivery of data is required, and there are many others—Web protocols like HTTP(S), VoIP protocols like SIP, and so on. From the perspective of the platform complementor, all these protocols are a part of the platform assets provided by the Internet—a complementor would say that they are "part of" the Internet. However, from the perspective of the abstract architecture described above, most of these other protocols are not part of or supported by the Internet, in that the routers that forward packets inside the Internet are not involved in the correct operation of these protocols. They are implemented in the end-nodes attached to the Internet. The software in the end-node is an important part of the platform assets on which application developers can depend. At the same time, since the public Internet defined by our abstract architecture is not involved in the operation of these protocols, application designers are free to ignore them in favor of alternatives specifically suited to their needs.

Second, content or application providers in different stages of development may adopt different strategies for the use of the public Internet and other platform assets. For example, Skype and Vonage both offer telephony (and videoconferencing, chat and other electronic communication services) that rely on their customers' broadband connections to originate calls. These Voice-over-Internet services compete directly with telephony services offered by legacy telephone and cable companies like Charter, Comcast, and Verizon; however, those latter providers use separate IP networks when their subscribers originate calls from their homes. To better compete with each other, Skype and Vonage allow their customers to make telephone calls to destination telephones that are not connected via broadband connections and many of the legacy telephone and cable

---

[26] The forwarding function, of course, depends on the computation of routing information, which is done by a suite of routing protocols and mechanisms in the Internet. However, these were not seen as a core part of the Internet's specification, and in fact have been replaced during the life of the Internet.

providers now offer applications that allow their subscribers to originate calls via broadband access connections should they elect to do so.[27] All of these providers are using Voice-over-IP (VoIP) technology and all of them are using a mix of Internet and non-Internet packet forwarding services to provide their customer's calling services.

The complementors are taking advantage of the generality and open access capabilities of the Internet and using those capabilities, potentially in conjunction with other non-Internet, potentially privately owned or customized capabilities, to provide services that are narrower in scope but more specialized to the needs and requirements of their target customers. Innovations in the Internet can result in changes in complementors' behavior and induce complementor innovations that then feedback to the Internet, driving the sorts of increasing returns to scale that characterize GPTs. For example, Netflix only began delivering content via the Internet when the Internet became capable of supporting broadband data rates; and as those rates increased, Netflix has increased the resolution of its videos. The innovation in data rates drives innovations in the content and applications that make use of the faster data rates, which in turn drives the need for faster hardware and software (and innovation by the complementors who provide the equipment on which the Internet runs).

Third, and perhaps most important for our purposes here, the role that the public Internet plays in creating the set of platform assets available to innovators is fluid and changeable. Further, the industry structure can be expected to be highly dynamic. The firms that today implement the public Internet (as characterized by its abstract architecture) may also offer other platform assets, and these may become woven into the service that users come to think of as the public Internet. A decade or so ago, it was relatively easy to identify the Internet Service Providers (ISP). The ISPs were (relatively speaking) a more homogeneous group and the range of services they provided were reasonably cleanly identifiable as data communication services. Today, with CDNs, cloud service providers, application and content providers all making use of and adding to the functionality of the Internet to varying degrees, it is much more difficult to craft a clear definition to decide which businesses or services are Internet services and which are not.

Although most of the focus here is on platform complementors—edge providers that offer applications and content "over the Internet," the firms that provide the basic hardware and software that allow the Internet to run are also network complementors. This includes firms like Cisco, Huawei, Microsoft, Oracle, and Intel.

These network complementors may regard the Internet as a source of demand for their goods and services, and consideration of this perspective is relevant to a layered-view of the Internet. In a layered perspective, the complementors exist at the layers above and below the Internet. In considering what is essential Internet functionality, due consideration has to be made for what the complementors need from the Internet layer and what the Internet layer can expect from the complementors.

---

[27] Enabling the applications allows customers to use their home service when away from home, something Vonage and Skype have always allowed.

Although it may not be precise or always clear who is a network complementor, consideration may be made to the range of services and business-to-business relationships that have evolved between firms engaged in the evolving Internet ecosystem. Our characterization of the public Internet through this lens is intended to be a more expansive conception than the abstract, architectural layering.

3.3. Customer Experience

Our third lens is at the highest level and focuses on the customer experience. This refers to the range of services, content and applications that end-users can expect to access via the Internet and their experiences when doing so. This includes the search capability, which is essential functionality for users to be able to make use of the Internet as a market. The search capability is core economic functionality that allows customers to identify what goods and services, including content, they can use the Internet to access. The choices of what consumers can do, how they make those choices, and what they experience when they make those choices play into the notion of the customer experience. As with the other layers, this is a moving target as the Internet continues to evolve.

For example, users expect to be able to browse the Web, use email and chat, participate in social network applications like Facebook and Twitter, engage in eCommerce via sites like Amazon and eBay, and stream media from YouTube and Netflix. Users expect to be able to have choice and control over the applications and services they use, and the devices they utilize to access the Internet. Users expect to be able to reach websites that are identified by their DNS names and, at least in principle, be able to establish communications with any IP address in the global name space (although in many cases, end-users will need to rely on the DNS or other services in the Internet to resolve their search queries into IP address destinations).

For these choices to be viable, users expect that their experiences when using these applications and services will be of sufficient quality to be usable, but end-users generally recognize that not all applications will be usable on all devices everywhere. For example, below some minimal data rate or in excess of some level of end-to-end latency, certain applications may either not work or provide such an unacceptably low level of performance as to be unusable. Real-time telephony requires that roundtrip delays be less than about 250 milliseconds or full-duplex conversations have to resort to a walkie-talkie mode of communication. Higher-resolution streaming video and fast-response interactive gaming may have both minimum data rate and maximum latency tolerances.

Some of the services that end-users may or may not[28] recognize as Internet services (like telephony and increasingly television) may be provided via the Internet, but are not uniquely, principally or historically identified as Internet services. Indeed, most of the Voice-over-IP or IPTV traffic is transported via specialized services or IP networks that are separate from or only partially overlap with the public Internet. Other services, like Web access, email, FTP, search, or DNS are principally, uniquely, and historically identified as Internet services. Some other services like text

---

[28] A Vonage customer may think of telephony now as an Internet service, while many other subscribers may not be aware of how the Internet or its supporting technologies (like IP) may be used to support telephony services.

messaging or chat are somewhere in-between, both having pre-Internet, mobile telephony precursors. Newer services like video conferencing, which also had pre-Internet precursors (e.g., ISDN and other specialized systems), have achieved true mass-market success only once they became easily accessible via the Internet.

Although consumers may not care about precise definitions, they are likely to care about who is responsible for meeting their expectations for quality and the terms of their service level agreement (SLA) for the Internet services they pay for. That certainly includes the ISP from whence they get their Internet access connection,[29] but it may also include edge providers of content, devices, applications, and ancillary services.[30] When problems arise, as much as consumers might like, there is often no single party responsible for resolving those problems.

For most customers, the technical details of how the experience is supported and what firms, networks, or capabilities have to be aligned and participate to make the experience feasible is either irrelevant or beyond their understanding. Section 3.2 described the Internet through the lens of the platform complementors. It made the point that in today's Internet there are a range of platform assets, not just the public Internet. It distinguished the set of actors that are responsible for creating those platform assets, including Internet service providers, providers of cloud and CDN services, and the like, and the set of actors that exploit those platform assets, the application developers. Through the lens of the user experience, none of this structure is visible. Nor should we expect the user to understand all these distinctions. The user sees the totality of the Internet ecosystem through the lens of his experience using it.

The consequence of this blurring of the components that make up the user experience becomes clear when something goes wrong. None of the layers of the Internet were expected to work perfectly all the time. At the packet forwarding layer, designers understood that routers could fail, links could fail, packets would be lost from time to time, and so on. For example, to deal with link and router failures, dynamic routing protocols were designed, and in general, each layer of the Internet attempts to deal with and mitigate the failures and limitations of the layers beneath it. For example, when a distributor of content decides to use a CDN for that purpose, the resulting user experience depends on fewer parts of the Internet (because the content is hosted at a point near the end user), so the experience is less likely to be impaired by network failures.

But when some problem does manifest at the level of the user experience, the user has no way of deconstructing the elements that created that experience and assigning the responsibility to the correct actor. This should not be the job of the user, but that raises the obvious question of where that responsibility should lie. One of the limitations of the current Internet ecosystem is that it is

---

[29] As we have discussed elsewhere (Bauer et al., 2016, 2015, and 2010), the end-user's experience regarding the performance of broadband Internet access services is not solely under the control of the access ISP. For example, many Internet access problems may be attributed to poorly performing customer WiFi networks.

[30] Internet access devices come in many flavors from PCs to smartphones to eReaders to purpose-built devices (IoT devices like NEST thermostats) and may come bundled with, or allow users to separately add, software applications enabling additional Internet capabilities (e.g., a browser, capability to access protected content/services). Consumers may pay and contract for specialized Internet services like Vonage (VoIP), Netflix, or access to the WSJ.

not well equipped with tools that can analyze problems (whether absolute failures or degradation of performance) and trigger mechanisms that can deal with the problem. But calls for "a better Internet", whether better means more secure, more resilience, higher performance, or whatever, are not actionable using only the lens of the user experience. Whatever the issue, the problem must indeed be deconstructed and the responsibility for the mitigation be assigned to the appropriate actors for any progress to happen. Those who hope to translate aspirations into action must understand the layering and modularity of the Internet ecosystem, and how the overall user experience is composed out of function and services provided by the different elements.

4. Using our lenses to Understand Internet Trends

To understand how the above characterizations may be useful, it is worth considering some specific examples.

For example, as a thought experiment, imagine that Facebook, an application platform that is currently implemented on top of the Internet (with a lot of non-Internet, private network services and infrastructure added in) were capable of delivering an end-user experience that was an essentially complete substitute for the end-user experience we associate with the Internet.

With the addition of email, chat, streaming video, news and other content and services that are typically associated with the Internet, the Facebook experience is becoming more Internet-like already, so the thought experiment is not purely hypothetical. If the user experience became a complete substitute for the Internet, how might we differentiate Facebook from the Internet? From an end-users' perspective, there might not be any reason to care; but as we will argue further below, there are some fundamental differences between such an evolutionary path that would have relevance to network engineers, economists, and Internet policymakers that are worth keeping clear.

By assumption the customer experience offered by the evolved Facebook was a complete substitute for an Internet experience, but we would still conclude that it was not Internet-like, in that its design and implementation would be controlled by a single company. The centralized control violates an important architectural feature, namely that control is distributed and decentralized. There is no reason to believe that this hypothetical future Facebook, even if it matched all of today's user experience, would be an open and effective platform for third-party innovation. It would fail the "platform" characterization even if it passed the "experience" characterization to some extent.

Alternatively, were someone to argue that the fact that many applications and services like Google that rely on CDNs that make use of non-public Internet resources, including private networks, renders those services as not valid parts of the Internet experience, we would disagree. Failing to include popular services that the general public clearly regard as associated with the Internet as part of the Internet experience would render the concept overly restrictive and arbitrary. What makes Google part of the Internet experience is how consumers access it and the use Google makes of basic Internet functionality to connect with end-users. Google is a complementor that builds on top of the Internet, and by so doing, participates in sustaining the Internet experience. However, Google's internal network is not-Internet-like because it is not a general purpose, open access platform. This highlights how network elements may be regarded as contributing to the consumer experience, and yet not be regarded as part of the Internet.

In the following sub-sections, we highlight some of the trends that are changing the Internet. We use our three-lens framework to explore the implications of these trends.

4.1. Changing topology of interconnection and routing

The legacy Internet was a network of hierarchically organized ISPs. At the top of the hierarchy were the Tier 1 ISPs that collectively defined the global address space. The Tier 1 ISPs interconnected directly with each other via revenue neutral peering agreements. Lower Tier ISPs purchased transit services from higher level ISPs in order to deliver traffic to addresses that were not in their domains. The ISPs offered only a single grade of best-effort packet delivery service.

Tier 1 ISPs only handed off traffic to other Tier 1 ISPs that they could not terminate directly. Lower Tier ISPs interconnected with higher level ISPs and purchased transit services to deliver traffic they could not deliver via themselves or via their peering agreements with higher-level ISPs. The legacy Internet offered a single grade of best-effort service that lacked quality-of-service guarantees.

This early picture of how the Internet is interconnected for end-to-end packet delivery has evolved greatly. A first cause for the flattening of this hierarchical structure is increased peering between networks.[31] The trend began with direct peering links between pairs of networks at colocation facilities and accelerated with the popularity of Internet Exchanges Points (IXPs), whereby a single network can easily peer with hundreds of other networks. Membership at IXPs is further fueled by the emergence of *remote peering*. Today, networks can connect to IXPs in geographically distant locations, leveraging connectivity provided by remote peering providers, another example of providers offering a specialized packet forwarding service that may be based on IP but is not necessarily part of the Internet.

For example, DE-CIX, one of the largest IXPs in the world by number of connected networks and traffic exchanged, offers remote peering availability (through partners, referred to as *resellers*) in more than 450 cities in more than 70 countries.[32] A study conducted in 2014 found significant portions of networks connecting to major IXPs do so remotely.[33] Remote peering represents an important architectural trend, since it flattens the AS-level topology, removing the dependence on intermediate transit networks. It also highlights changing interconnection economics, since remote peering often provides a less expensive alternative when compared to legacy transit services. Remote peering providers (and, more recently IXPs themselves) present us with an additional network complementor, offering wholesale wide-area connectivity to networks all over the globe that wish to participate at chosen IXPs. While IXPs have provided the dominant mode for

---

[31] For further discussion of how interconnection has evolved in the Internet from a binary world of transit and peering to the more complex world that also includes paid peering, partial transit and other hybrid interconnection models, see Clark et al. (2006, 2011).

[32] https://www.de-cix.net/en/access/where-to-connect.

[33] Castro et al. (2014).

interconnection in Europe for a while, in recent years, they have been gaining ground worldwide with 640 IXPs operating currently in 128 countries.[34]

A second significant cause for the flattening of the AS-level topology is the growth of CDN's and major content providers themselves deploying their servers in access networks.[35] The dramatic changes in the patterns of Internet traffic first came to widespread attention with the publication of the Arbor Networks study of traffic patterns in 2009[36] that documented the rise of the "hyper-giants." This showed that a few very large content providers (like Amazon, Google, and Netflix) were connecting directly to the so-called "eyeball" networks that provided content directly to their access subscribers. In 2007, thousands of ASes accounted for 50% of the traffic delivered to end-users; whereas by 2009, only 150 ASes contributed 50% of all the traffic. These trends appear to be continuing. Today, access ISPs see a much larger share of their traffic coming from an even smaller number of much bigger content providers, and aggregate traffic volumes continue to grow exponentially. Much of this growth is being driven by video (and most of that is entertainment-related).[37]

Hyper-giants increasingly rely on their own fiber links to interconnect data centers, thus bypassing transit providers and delivering content directly to respective eyeball networks. Additionally, more and more content providers deploy their own content caches either directly within eyeball networks or at IXPs or colocation facilities, effectively operating their own CDNs.[38]

A byproduct of these developments is that ISPs who depended on transit revenues are seeing that source of funding wither. As Huston (2017) has pointed out, transit revenues provided an important revenue source that sustained investments in backbone Internet capacity. With that funding source disappearing, ISPs without a direct revenue interface with end-customers may find themselves starved for funds to continue investing in backbone facilities. Is that a problem? It might be, but then again maybe not.

Although revenue from transit traffic may wither, demand for dark fiber in the U.S. and worldwide continues to grow rapidly[39] and we even observe a recent upswing in submarine cable construction worldwide.[40] Thus, the decreasing reliance on transit service has not yet resulted in fewer investments into backbone infrastructure (i.e., fiber networks). While bandwidth growth rates of

---

[34] This estimate is from the website www.peeringdb.com which listed over 640 IXP's located across 128 countries, including 240 IXP's in Europe, 141 in North America, 107 in Asia, 75 in South America, 46 in Africa, and 32 in Oceania.

[35] See Stocker et al. (2017).

[36] See Labovitz et al. (2009).

[37] See Lehr & Sicker (2017).

[38] See Böttger et al. (2018) for a recent study that describes how Netflix's in-house CDN exploits the benefits of IXPs, while also deploying Netflix caches within eyeball access networks.

[39] See https://globenewswire.com/news-release/2018/03/15/1437953/0/en/Dark-Fiber-Networks-Market-is-Expected-to-Hit-US-11-Billion-By-2026-Credence-Research.html.

[40] See https://www2.telegeography.com/new-submarine-cable-builds.

classic transit ISPs might be on the decline, enterprises, hyper-giants, cloud providers, as well as providers of customized network services (e.g., the remote peering service providers discussed above) have been expanding their network capacity in response to growing demand for infrastructure.

4.2. Rise of Content/Access Censorship, Filtering, and Limited Access

Key features of the user experience are that it enables users a degree of control and freedom of choice to access the content and applications they want, and to interconnect the devices they want to connect, subject to the proviso that the content, applications, and devices do not cause harm to the Internet or its end-users. This freedom of choice was embedded in FCC Chairman Powell's four policy principles that he articulated for the first time in 2004.[41] The principle of freedom of choice is consistent with the Internet viewed through the abstract architectural lens and thus overt efforts to censor or limit Internet users' choices would be inconsistent with our characterization of the Internet. Moreover, from the perspective of the customer experience, the desire that the Internet experience be safe, or at least not unreasonably unsafe, seems a reasonable user expectation and seeking to promote that goal a reasonable policy objective. Nevertheless, the abstract architectural lens does not provide guidance on how that might be ensured and what sorts of limitations (in terms of access to content, applications, or devices) might be appropriate to promote safety. Thus, we might expect complementors (which might include end-users taking responsibility for their own safety) to add (or limit) capabilities with the goal of better promoting a safe experience.

Consistent with this, we see makers of tablets for children explicitly limiting the capabilities of those tablets to prevent them from making full use of the Internet's capabilities (e.g., by including parental controls and hand-curated browsers and content libraries to allow parents to censor what their children are able to access via the Internet).[42] We do not view such efforts as inconsistent with the Internet concept since the complementors are expanding the range of choice available to Internet end-users (in this case by selectively enabling Internet functionality and by augmenting it with curated content, services, and applications). Such efforts represent an enhancement to the end-user experience. While most people would accept the rights of parents to censor what content and applications their children are allowed to watch, and many may also appreciate efforts by ISPs or providers of other adjunct services to filter out malware and other harmful traffic, people may

---

[41] Powell's four principles stated that consumers should have freedom of choice to (1) access legal content, (2) use lawful applications, (3) connect safe devices, and (4) select among a competitive selection of choices for service, application and content providers (see Remarks of Michael K. Powell, Preserving Internet Freedom: Guiding Principles for Industry, prepared for Silicon Flatirons Symposium, Boulder CO, February 8, 2004, available at http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-243556A1.pdf; and FCC (2005) *Policy Statement*, In the matter of appropriate framework for Broadband Access to the Internet over Wireline Facilities, CC Docket No. 02-33 (and related matters/dockets), adopted August 5, 2005, available at: http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-05-151A1.pdf).

[42] A number of tablet vendors, including Amazon and Kurio make tablets targeted at kids. In addition to being less expensive and easier to use, these also come with range of tools for limiting what services, content, and websites may be accessed. A common add-on is a "safe browser." For a sample review of such devices, see "The best kids' tablets for 2018," Tech Advisor, available at https://www.techadvisor.co.uk/test-centre/tablets/best-kids-tablets-2018-3378548/ (visited July 9, 2018).

reasonably differ in their attitudes toward what they regard as arbitrary censorship or limitations on the Internet experience.

In a number of countries, Facebook has promoted its Free Basic service as a stripped down version of Internet access. Policymakers around the globe have argued about whether such subsidized but restricted access promotes greater Internet accessibility and so contributes to the growth of the Internet or breaks a fundamental promise of the Internet as an open access platform. Critics of Free Basic have argued that the sort of restricted access it provides is inconsistent with core principles of the Internet since it interferes with end-users ability to access any content or application they wish.

To judge whether a service such as Free Basics is consistent with the core principles of the Internet, one can look at such a service through our three lenses. In terms of user experience, a key question is whether a user of Free Basics is locked into that service if the user subscribes to it, or whether the user can fluidly step outside the service and partake of other parts of the experience (for the normal usage fee) as desired. A version that locked a user into a "walled garden" would be inconsistent with the characterization of the Internet as user experience; however, an experience that allows full access, but with certain parts of the experience available for a lower cost does not seem inconsistent. Through the lens of the platform complementor, the test would be whether the option of participating in the Free Basics program is open to all application providers ("permission-less innovation") or requires approval from a controller of the platform. Through the lens of the abstract architecture, the test would be whether subscription to the Free Basics service prevents a user from the open exchange of packets with arbitrary end-points (again, for a normal usage fee). Different variants of a Free Basics service would be judged more or less conformant with our criteria for being a part of the Internet.

In another context, some analysts point to the filtering and blocking that the Chinese government undertakes to manage access to information by its citizens in China via the Internet as proof that China is cutting itself (its citizens) off from the Internet. However, it is possible to reach end points in China and establish basic connectivity. Through the lens of abstract architecture, the Chinese Internet is indeed a part of the global public Internet. Additionally, the Internet in China provides Chinese application complementors with similar functionality to what applications providers use elsewhere. The most significant distortion of the Internet experience in China, relative to the rest of the world, is probably that China blocks the use of applications such as Facebook. The dominant application that most consumers use in China to access the Internet is a mobile application called WeChat, which provides a full-spectrum of Internet services like chat, messaging, web browsing, et cetera.[43] This would seem to offer a reasonable Internet customer experience. However, many proponents of the Internet's role in supporting free speech and open access are concerned by the heavy-handed regulation that the Chinese government imposes on Chinese ISPs and complementors. One might argue that the Chinese government exerts excessive centralized control

---

[43] WeChat is an application platform provided by the Chinese high-tech firm TenCent (see "Tencent, the $500Bn Chinese tech firm you may never have heard of," The Guardian, 13 January 2018, available at https://www.theguardian.com/business/2018/jan/13/tencent-the-500bn-chinese-tech-firm-you-may-never-have-heard-of). TenCent's focus so far has been on addressing the Chinese domestic market, which requires it to be sensitive to the interests of Chinese regulators; however, one might make a similar claim about Facebook and its need to be sensitive to political pressures in the United States.

over WeChat, and as such poses an undue risk to a level of content control or censorship that violates fundamental principles of Internet openness and an abrogation of the potential for permission-less innovation.[44]

What this discussion highlights is the complexity of trying to constrain too narrowly how we think about the Internet in terms of the topology of traffic flows and the changing balances of power in who gets to decide what traffic flows where and how. For many, the changing Internet that has moved us very far from the original model of a hierarchy of interconnected ISPs to a much more complex ecosystem with hyper-giants and overlays has enabled much more capable and richer customer experiences. Many of the enhancements (e.g., better support for QoS, improved reliability, expanded access to content) are presumptively consistent with improvements to the Internet that everyone applauds. On the other hand, as the Internet has become more embedded in our public discourse and plays a bigger role in shaping public opinion, the Internet is justifiably attracting more attention as a public forum for citizen discourse, for its role in international relations, and for its impact on sovereign power.

4.3. Emergence of clouds

As we discussed in section 3.2, the platform complementors today have a rich suite of platform assets from which to draw in designing their applications. The concept of "cloud computing" is that storage and computing capabilities are available (from large data centers) that can be requisitioned on demand by applications. Applications today can be developed without any major investment in computing and storage hardware because computing and storage can be purchased from a competitive set of cloud providers. Cloud is another example of a general purpose technology (GPT) that can be used in a number of ways to generate downstream value.

Cloud, as described above, is not a part of the public Internet as characterized through the abstract architecture lens. The data centers that provide cloud services are just another sort of end point device (albeit a very complex and specialized device) attached to the Internet. More recently, major cloud providers have started offering enterprises the possibility to lease dedicated connections (separate from the public Internet) to their cloud platforms, offering enhanced performance for bandwidth-intense and latency-sensitive applications. In such cases, cloud services can be used without traversing or relying on the public Internet at all.[45] But cloud computing is clearly part of the Internet ecosystem, specifically, it is a platform for application development.

In the future, computing and storage capabilities may actually move "into" the public Internet. They may be deployed in 5G wireless base stations, for example. Ongoing research on alternative

---

[44] As noted in the footnote 43, the need for Internet and complementor firms to comply with regulations and be sensitive to political interests is hardly unique to China. Perhaps all that can be said is that there are significant differences in the extent of government control over the Internet that may be excessive in some countries. The notion that a government may have a direct equity stake in a firm with a leading role in shaping the Internet experience for a significant share of the world's population may be viewed by many as asserting an excessive degree of control (see Yuan, L. (2017), "Beijing Pushes for a Direct Hand in China's Big Tech Firms," Wall Street Journal Online, 11 October 2017).

[45] See Microsoft's "ExpressRoute" offerings (https://docs.microsoft.com/en-us/azure/expressroute/) or Amazons "AWS Direct Connect" offerings (https://aws.amazon.com/directconnect/details/).

future architectures for the Internet clearly imagine that these services might be "in" the network—integrated into the packet forwarding function.[46] A move to add in-network storage and computing capabilities to the portfolio of network services would challenge the current concept of the public Internet; it would add new in-network functionality, potentially with a level of generalness that is consistent with the generality with which the Internet supported transport when that was the Internet's principal function.

From a customer experience perspective, the emergence of cloud computing is just another buzz-word that the user hears to describe the current Internet ecosystem. Today, we view cloud service providers as implementing part of the platform assets of the Internet, but not all such providers would be viewed the same. From a complementor perspective, we would view cloud service providers as platform providers if they provide general purpose computing and storage functionality that is consistent with the Internet's character of providing a general purpose platform open to all users, recognizing that in the future, such capabilities may be embedded and part of the basic functionality supported in the Internet (rather than dependent on third-party providers). At the same time, we would view more specialized cloud service providers as customers or on-top complementors. Where to draw the line may not always be clear and may change over time as technology and markets continue to evolve.

4.4. Virtualization

We see today that raw packet forwarding capacity can be sliced up into shares, one share of which might implement a part of the public Internet (as defined by its abstract architecture) and others of which might be used to build other IP-based networks, which either are used internally to one firm, or are offered as platform assets to other developers. Cloud computing has the same character—a large physical computing and storage facility is sliced or divided up into shares that can be used by different developers. The ability to create these shares is becoming more dynamic, with cloud computing available essentially instantaneously on demand. The end-point of this trajectory might be a world in which an application developer can dynamically requisition a complete network, computing, and storage capability as needed. Once a developer can get any ICT resource as a virtual slice, on-demand and can rent arbitrary bundles (from raw CPU cycles to turn-key final product solutions) it is not clear what the concept of a "public Internet" means. We are still a long way from full realization the vision of virtualization, but it does represent a trend to watch.

4.5. Reduction in global addressability

Another trend is the increasing cleavage of the once globally unique IPv4 address space into fragmented ISP-internal networks. This trend does not only manifest in internal networks of content providers or clouds (as discussed above), but also increasingly within eyeball networks, where ISPs assign internal IP address space to their subscribers and only connect them via large-scale NAT gateways (Carrier-Grade NATs or "CGNs") to the public Internet. In eyeball networks,

---

[46] Proposals for alternative Internet designs with in-network computing or storage include Nebula (see Anderson et al., 2014), Mobility First (see http://mobilityfirst.winlab.rutgers.edu/) and Named Data Networking (see Zhang et al., 2014).

this trend is primarily driven by the exhaustion of the globally routable IPv4 address space.[47] While NATs have long been employed within home networks (i.e., as part of Customer Premises Equipment, home routers), ISPs are increasingly pushed to deploy CGNs at scale, handing out internal IPv4 addresses to their subscribers and translating them at gateways within the ISP's network. The situation is particularly severe in cellular networks, where some 95% of ISPs already deploy CGN.[48] This implies that the majority of Internet users as of today access the Internet via CGN gateways.

NATs directly break two of the core architectural principles outlined earlier: end-to-end connectivity and a globally unique address space. This has consequences for innovation, since it makes development of peer-to-peer based applications difficult, and in many instances impossible. While techniques exist to circumvent the limitations imposed by well-known home NAT devices have long existed,[49] they often cannot be applied in the case of CGNs. Moreover, because of the necessary multiplexing of end-user's connections onto fewer public IPv4 addresses, CGNs necessarily impose limits and quotas on the number of possible concurrent connections per subscriber. IPv6 is supposed to replace IPv4 as the dominant addressing protocol in the Internet, effectively solving IP address scarcity issues. Yet, as of 2018 the majority of end users worldwide still do not have IPv6 connectivity, and the majority of end users have only NATed IPv4 connectivity, since most users connect via cellular operators.[50]

Relaying this back to our question of what constitutes a legitimate form of Internet access, it is worth noting that there are currently no regulations regarding IPv4 CGN or IPv6 deployment. Does NATed IPv4 connectivity with limited connection quotas and the impossibility to use certain peer-to-peer applications count as a legitimate form of Internet access? Could IPv6 connectivity become a necessity for Internet access to qualify as 'legitimate'? These are questions that we expect the Internet community to need to confront in the near term.

4.6. Internet of Things

The term Internet of Things (IoT) refers to devices (often small and with limited capabilities) that perform sensing and control without direct human oversight or supervision. Some of these devices (such as surveillance cameras) are designed to be used by people, others (like thermostats) may perform their functions automatically. Some of these devices are attached to the Internet in a fully general way, even though they are fixed function devices at the application layer. "Things" differ from more typical Internet devices in that they are not general purpose platforms onto which

---

[47] For more details, see Richter et al. (2015).

[48] For more details, see Richter et al. (2016).

[49] For example, UDP can be used for "hole-punching" to create a direct connection between an end-node inside the firewall of a NAT (see https://en.wikipedia.org/wiki/Hole_punching_(networking)). Also, see for example, "STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)," RFC 3489, March 2003, available at https://www.ietf.org/rfc/rfc3489.txt.

[50] According to Vint Cerf, speaking on behalf of Google, as of January 2018, "roughly a quarter of users access Google over IPv6" (see https://www.zdnet.com/article/googles-vint-cerf-quarter-of-internet-is-ipv6-but-heres-why-thats-not-enough/).

applications are later loaded, but instead are provided as an integrated hardware and application solution.

While some IoT devices are directly connected to the Internet in a way typical of a workstation or mobile device, others take advantage of specialized network technology that is low power, inexpensive to integrate into the device, and suited only for the task at hand. The question arises as to whether these devices that rely on specialized (and more limited) interfaces are a part of the Internet. The answer depends on which lens we use to answer this question. Using our abstract architecture, these devices are not directly a part of the Internet. The network technology that carries their traffic is not designed to provide general, application independent packet carriage to other end-points on the Internet. The interconnection of these specialized networks to the public Internet (if that connection happens at all) may be through a device that implements application-specific traffic forwarding, not general, application-independent forwarding. On the other hand, these devices may embody an application that the user considers to be part of the Internet experience. The application developer has taken advantage of these specialized networks as one of the potential IoT platform assets and woven these assets together in the context of the specific application to create a new sort of Internet experience. For many of these complementors, end-users may not even view the complementors as part of the Internet ecosystem. For example, a washing machine or refrigerator may be connected to the Internet in multiple ways – using either a standard Internet connection (e.g., potentially controlled via a user-accessible smartphone application that may communicate via the home owner's standard broadband connection) or via one of the specialized sorts of connections described above. In the former case, we would view the appliance as being on the Internet in the same way as a PC may be on the Internet; while in the latter we would require a more nuanced answer as noted above.

4.7. Entertainment video dominates traffic

According to Cisco's VNI forecast, video will account for 82% of global IP traffic by 2021, up from 73 percent in 2016.[51] A significant driver for the rise of the hyper-giants and the flattening topology of the Internet discussed earlier is associated with the growth of over-the-top video from Netflix, Google's YouTube, and other video services. Most of this video traffic is entertainment media (movies, television programming, and short-form videos). Also, it is worth noting that even if IP video already constitutes the majority of IP traffic, shifting all of today's entertainment media traffic to the Internet would require massive expansion of Internet capacity, and further re-architecting of the Internet (e.g., even more hosting of content close to the edge in access networks).[52] At the same time, many of the larger access ISPs have been investing billions of dollars in the acquisition of media content properties (e.g., AT&T's acquired Time Warner for $85

---

[51] See https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/vni-hyperconnectivity-wp.html (visited August 14, 2018).

[52] Most of the entertainment video traffic is currently carried via over-the-air, direct broadcast satellite, cable television or other wired networks that are separate from the Internet, and in most cases, still not IP, although that is changing.

billion in 2018; Verizon acquired AOL for $4.4 billion in 2015 and Yahoo! for $4.5 billion in 2017; and Comcast acquired NBC for 430 billion in 2011.[53]

In Lehr & Sicker (2017), we examined the implications of the growth of entertainment video and the potential implications that might have for the architecture, business models, and regulatory policies for broadband platforms and the Internet. We pointed out that the fundamental economics of the entertainment industries are distinct and different from many of the other types of uses for which we view the Internet as valuable. For example, the industry value chains and economics for healthcare, the distribution of electric power, and the automation of transport are quite different from those of entertainment and media industries. We posited that even if the majority of video is delivered via IP networks in the future and makes use of broadband access platforms, that need not mean that the IP video would necessarily be carried over the broadband Internet access connection and, in principle, at least, could largely be carried via a separate IP networks from the Internet.

From our abstract architectural lens, this growth of entertainment-focused video traffic does not obviously violate any of the essential features we identified in our Table 1. The bulk of today's entertainment media traffic has certain characteristics that may distinguish it from other types of application traffic, including the fact that most of it is one-way (downstream) and cacheable, which distinguishes it from certain other types of traffic such as video-conferencing (which is two-way and not cacheable). However, our view of the Internet does not require that traffic fit any particular pattern. Furthermore, entertainment media is evolving to include interactivity and other networking capabilities that require the sorts of flexible support needed by other demanding types of applications. Thus, there is no a priori reason to believe in advance that an Internet that evolved to support mostly entertainment traffic would necessarily be less capable or less able to support other diverse applications (like healthcare or smart automation applications).

From the perspective of network complementors, the entertainment industry has a long history of wanting to bundle content with distribution to enable tight control over the end-user experience in order to enable differentiated pricing. "Windowing" or phased delivery to effect price discrimination has a long history in entertainment media. In an earlier age, this took the form of first releasing books as high-priced hardcovers or movies to theaters to capture the high-pay-per-view audience and later distributing to other lower-price-per-view channels (paper-backs, television serialization, DVD discs) to extract revenues from consumers. With the rise of more complex media distribution options, including the Internet, the "windowing" has become more complex and targeted, but the basic idea has remained: the owners of the content seek to control the modes of distribution to effect price discrimination to recover revenues. These strategies are necessary, in part, because most of the costs of content creation are incurred up-front and are sunk when it is time to consume (view) the content. Copyright laws provide a legal property right that may be asserted by content rights owners to allow them to recover the economic costs of creating

---

[53] See Diane Bartz & David Shepardson, *AT&T closes $85 Billion deal for Time Warner*, Reuters Business News, June 14, 2018, https://reut.rs/2K0RJea, Hamza Shaban, *It's official: Verizon finally buys Yahoo*, Washington Post, June 13, 2017, https://wapo.st/2NDHBqb, and David Gelles, *Verizon Bets on Video Ads in $4 Billion Deal for AOL*, New York Times, May 12, 2015, https://nyti.ms/2JKwTvr; and, Sam Gustin, *Comcast's NBCUniversal Deal: as One Media Era Ends, Another Begins*, Time Magazine, February 14, 2013, https://ti.me/2MirYYL.

the content, and thereby, contributes to ensuring adequate incentives for content producers to invest in content creation. Additionally, while there are confidentiality, security, and privacy concerns that arise associated with entertainment media, these may be of a different level of importance when compared to other types of application traffic (e.g., system control information or healthcare data). Finally, we note that there are likely to be systematic cultural differences (e.g., attitudes toward pornography, choice of native language, or preference for domestically originated content[54]) that might be more acceptably justified as amenable to national content-related regulations than for other types of Internet traffic.

While there is no a priori reason to conclude that the networking requirements of entertainment media are systematically different in some important way from other types of traffic that we would like the Internet to be able to support, it is not unreasonable to suspect that such differences *may arise*. Consequently, it is worth at least asking whether an Internet that is optimized to address the needs of the entertainment media industry is the Internet we want to protect and preserve, and to the extent we conclude that undesirable biases are being introduced (e.g., in terms of what capabilities are supported, the Internet's flexible support for diverse traffic types, asymmetric capacity provisioning, or excessive acceptance of traffic filtering), then we may wish to consider how policymakers might best act to off-set those adverse effects. One approach we suggested in Lehr & Sicker (2017) was to consider shifting a significant share of the entertainment traffic to another IP network (i.e., not the Internet).

Finally, from the customer experience perspective, we recognize that a growing number of end-users are coming to expect to be able to access their entertainment media on all of their Internet devices, wherever and whenever they happen to be. Providers like Comcast and Verizon are offering services to allow their video subscribers to access their video content via broadband access as part of their effort to hold onto their subscribers in the face of increased competition from over-the-top services.[55] Although end-users may not care or be aware of the network that is used to deliver the entertainment media to their device and even if most of the media content may not be delivered via the Internet, it is clearly the case that in some (maybe most) situations, the end-user experience will need to include the ability to access media content via the Internet. Moreover, it is not clear how one might distinguish between traffic that constitutes entertainment media and other types of Internet traffic, especially since much of this other traffic may also include video (e.g., from observation cameras, video conferencing, etcetera).

The growth of entertainment traffic and the consumer demand that is driving it is providing the revenue justification for much of the investment in expanding the capacity and capabilities of the Internet, which is generally perceived as beneficial for all potential uses. Therefore, although we do not think the growth of entertainment traffic necessarily poses a threat for the future evolution of the Internet, we recognize it as a trend that is worth watching since the dominance of any particular type of traffic or application might lead to the specialization of the Internet in a way that

---

[54] For example, Canada and France have strong domestic content rules for media content.

[55] Most subscribers to fixed broadband Internet access services purchase one or more additional services from their broadband Internet access service, usually either video television and/or voice telephony services.

would violate the generality and flexibility that has characterized the Internet to date and that we recognize as valuable features that are worth preserving.

4.8. Growing Concern with Cybersecurity

As the Internet provides a platform for merging the virtual and real worlds, embedding ICTs more deeply into the fabric of society and economy, more of what happens that matters will happen online. Criminal activity and the potential to suffer harm will go where the action is. Defending against cybercrime and other malicious activities is a requirement that is only going to get more important in the future.

Viewed through the lens of the user experience, the aspiration of better security is easy to state. However, through that lens, it is difficult to translate that aspiration into actionable tasks. The layered model of the Internet ecosystem can provide some structure to make better security actionable. Some aspects of security belong in the layer we call the public Internet as defined by abstract architecture. Making the global routing system (BGP) less vulnerable to malicious manipulation is a task that belongs at that level. As well, improving the security of some key supporting services (the Domain Name System and the Certificate Authority system) probably should be assigned to that layer of the system. However, many of the security problems the user faces today (spam, phishing, identity theft, malware downloads and the like) do not arise within the scope of the abstract architecture of the public Internet. They arise through design decisions of the platform complementors that make design choices at the application layer that create vulnerabilities and risky modes of operation.

It would be nice if the responsibility of improving the security of the Internet could be assigned to one set of actors (for example, the Internet Service Providers). However, that approach is not realistic and would be unlikely to succeed. It is necessary to have a more complex model of what is meant by "the Internet", and assign responsibility based on a taxonomy of the security problem. The scope of problems that can be addressed within the abstract architecture conception of the public Internet is in fact a small part of what, through the lens of the user experience, constitutes poor Internet security.[56]

5. Three Lenses and Questions for the Future

The idea for this paper emerged from our inability to come up with a single satisfactory answer to the question "what is the public Internet?" This leads us to ask whether adding "public" or capitalizing the "I" was important and in what context might the difference matter. We quickly realized that any single definition would be unsatisfactory and the whole exercise of providing a definition outside of the narrow context in which a definition might be necessary and used (e.g., in the language of a contract or a law, both cases where ambiguity can be problematic) was not very interesting. Yet, we recognize that experts and the general public often use "the Internet" in casual conversation, the trade press, academic research, and in policy debates as if everyone understands what those words mean. In light of the massive changes that have occurred in the economy and society, due in part to the global transition to a digital economy that the Internet has

---

[56] An extensive discussion of how to break the problem of Internet security into actionable parts can be found in Chapter 10 of Clark (2018), forthcoming.

played a significant role in, we think more clarity about what the concept "the Internet" ought to mean is valuable.

To address this challenge, we have proposed a three-lens framework to use when trying to determine whether some behavior, market development, or technology is consistent with "the Internet" or not. Ultimately, our hope is that this will allow us to focus on preserving the features of the Internet that we believe are worthwhile, while allowing the Internet to evolve. This will require us to avoid being confused or distracted by irrelevant changes, while focusing our efforts on addressing threats and opportunities that are essential to preserving the fundamental positive features of the Internet that we value.

The lowest level, which we refer to as abstract architectural, represents our minimalist characterization of the functionality or design principles that we think have been central to the core concept of what the Internet is and should be. While we may have more than one global communications and computing network infrastructure, we believe having at least one that has the features we identify in Table 1 is desirable and worth preserving.

Our middle lens focuses on network complementors, the enterprises that use the Internet as a platform to develop the applications and content that makes the Internet valuable and facilitates the Internet's role in the economy as a platform for innovation. In applying this lens, we ask what aspects of the Internet are necessary for complementors. In many cases, complementors make use of other platform assets that are not part of the Internet, but the end-result of what the complementors actions give rise to contributes to creating the customer experience.

The highest lens we apply is the customer experience. The focus here is on what we think are and ought to be reasonable expectations and aspirations for the Internet. Before broadband and better support for bounded latency transport existed, voice telephony and watching video over the Internet were not commercially viable. Today, end-users expect to be able to run such applications and those applications are now legitimately regarded as part of the customer experience. We believe that the performance capabilities of the Internet should continue to improve along multiple dimensions, including further reductions in latency, faster data rates, and additional functionality so that the Internet will not be the limiting factor in our abilities to offer better choices and higher quality services. We also think it is reasonable for consumers to expect that the customer experience on the Internet will be sufficiently safe that the users are not inhibited from partaking of that experience.

Enabling these enhancements and evolutionary (sometimes revolutionary) improvements often falls to the network complementors and often involves adding platform assets that are not part of the Internet. Figuring out when the capabilities delivered by those non-Internet assets ought to be considered as part of the critical infrastructure we associate with the Internet is an important question that engages network researchers and the enterprises engaged in providing the different components that support the Internet. It is an on-going challenge. Moreover, even after previously unavailable functionality becomes generally accessible in the Internet, we do not expect that all complementors will elect to make use of it. Indeed, requiring complementors to choose to use Internet functionality instead of private or non-Internet alternatives would be counter to the abstract architectural features in Table 1.

However, it is also possible that a powerful complementor (whether making use of the Internet as a platform for developing applications or content or as a provider of the equipment or technologies on which the Internet runs) might seek to threaten the Internet's ability to provide the alternative functionality (e.g., to foreclose competition from competitors that may choose to rely on the Internet). Any such threat should be watched with concern. A range of policy responses might be appropriate in response, ranging from antitrust to public subsidization for the threatened Internet functionality. The goal of this paper is not to recommend specific policy interventions, but rather to establish a basis for asserting a public interest so that an intervention might be justified if an appropriate one could be identified.

In the preceding section, we discussed a number of the most important trends that are characterizing the evolution of the Internet and what those may mean for how we might want to alter our core conception of the Internet or raise our awareness of threats, risks or opportunities that we see confronting the Internet.

When we apply our three lenses to today's Internet, we are generally optimistic. We find that today's Internet still fits lenses that we believe were also applicable to the legacy Internet. However, we also find potential concerns for the future that we think our three-lens framework helps elucidate.

For example, the reduction in global addressability makes it more difficult for the sorts of edge innovation that was implemented using peer-to-peer applications. We are not ready to conclude that this is a critical departure from the principles on which the Internet was built that may need to be rectified to preserve the Internet, but we do think it is something that bears watching. Another example is the rise of ever-more-capable platforms from complementors that may increasingly be viewed as substitute platforms for consumers, and even, application developers. The growth of Google Apps and Facebook in the U.S. and WeChat in China raise the concern that the generality and distributed control of the Internet might be lost in a future where a single enterprise effectively controls the user experience.

As long as the Internet continues to exist as an alternative source of the general purpose platform assets needed by complementors and there is sufficient competition among complementors, the rise of such powerful complementors in the market is not inconsistent with our view of the Internet – but again, it does raise concerns. Should any actor in the Internet ecosystem acquire excess market power, competition authorities should be concerned, and if any such actor were to take actions that threaten the core value propositions of the Internet, it may warrant policy action. Precisely what that action should be obviously depends on the context. A goal of our 3-lens framework is to help frame a policy perspective that is sufficiently flexible and relaxed as to avoid unnecessary public interventions in markets, but also capable of providing a basis for establishing a justification for public intervention on Internet policy grounds if needed.

6. References

1. Anderson, T., K. Birman, R. Broberg, M. Caesar, D. Comer, C. Cotton, M. Freedman, A. Haeberlen, Z. Ives, A. Krishnamurthy, and W. Lehr (2014), "A brief overview of the

NEBULA future internet architecture," *ACM SIGCOMM Computer Communication Review*, *44*(3), pp.81-86.

2. Bauer, S., W. Lehr and M. Mou (2016), "Improving the measurement and analysis of Gigabit Broadband Networks," TPRC44, September 2016, Alexandria, VA, available at SSRN: https://ssrn.com/abstract=2757050 or http://dx.doi.org/10.2139/ssrn.2757050.

3. Bauer, Steve, William Lehr, and Shirley Hung (2015), "Gigabit Broadband, Interconnection Propositions, and the Challenge of Managing Expectations," TPRC2015, Alexandria, VA, September 2015, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2586805

4. Bauer, S., D. Clark, and W. Lehr (2010), "Understanding Broadband Speed Measurements," (August 15, 2010). 38th Research Conference on Communication, Information, and Internet Policy (TPRC.org), October 2010, Arlington, VA, available at SSRN: http://ssrn.com/abstract=1988332.

5. Böttger, T., F. Cuadrado, G. Tyson, I. Castro, and S. Uhlig (2018), "Open Connect Everywhere: A Glimpse at the Internet Ecosystem Through the Lens of the Netflix CDN," in ACM SIGCOMM Computer Communications Review (CCR), 48(1).

6. Bresnahan, T. and M. Trajtenberg (1995), "General purpose technologies 'Engines of growth'?" *Journal of Econometrics, 65*(1), 83-108

7. Cannon, R. (2003), "Will the Real Internet Please Stand Up: An Attorney's Quest to Define the Internet," in Lorrie Cranor and Steven Wildman (eds) Rethinking Rights and Regulations: Institutional Responses to New Communications Technologies, MIT Press: Cambridge, MA, 2003 (available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=516603)

8. Castro, I., J. Cardona, S. Gorinsky, and P. Francois (2014) "Remote Peering: More Peering without Internet Flattening" Proceedings of ACM Conference on emerging Networking EXperiments and Technologies (CoNEXT)

9. Claffy, kc and D. Clark (2014), "Platform models for sustainable Internet regulation," *Journal of Information Policy*, *4*, pp.463-488, available at http://www.jstor.org/stable/10.5325/jinfopoli.4.2014.0463

10. Clark, D. (2018), Designing an Internet, MIT Press: Cambridge, MA, 2018 (forthcoming).

11. Clark, D. and kc Claffy (2015), "An Inventory of Aspirations for the Internet's future," Technical report, Center for Applied Internet Data Analysis (CAIDA), available at https://www.caida.org/publications/papers/2015/inventory_aspirations_internets_future/inventory_aspirations_internets_future.pdf

12. Clark, D., W. Lehr and S. Bauer (2011), "Interconnection in the Internet: the policy challenge," *39th Research Conference on Communications, Information and Internet Policy* (www.tprcweb.com), Alexandria, VA, September 2011, available at https://ssrn.com/abstract=1992641

13. Clark, D., W. Lehr, S. Bauer, P. Faratin, R. Sami, and J. Wroclawski (2006), "Overlay Networks and Future of the Internet," *Communications and Strategies*, no. 63 (3rd Quarter 2006) 1-21, available at ftp://ftp-sop.inria.fr/members/Didier.Parigot/PDFarticle/francesco/Clark2006.pdf.

14. Clark, D., J. Wroclawski, K. Sollins, and R. Braden, R. (2002), "Tussle in cyberspace: defining tomorrow's internet," *SIGCOMM Comput. Commun. Rev., 32*(4), 347-356.

15. Clarke, G., C. Qiang and L. Xu (2015), "The Internet as a general-purpose technology: Firm-level evidence from around the world," Policy Research Working Paper #7192, The World Bank, February 2015

16. CSTB (1994) *Realizing the Information Future: The Internet and Beyond*, Computer Science and Telcommunications Board (CSTB) of the National Research Council (NRC), Washington, DC: The National Academies Press (NAP), 1994.

17. De Rosnay, M. and H. Le Crosnier (2012), "An introduction to the digital commons: From common-pool resources to community governance," in *Building Institutions for Sustainable Scientific, Cultural and genetic Resources Commons*, International Association for the Study of the Commons, available at https://halshs.archives-ouvertes.fr/halshs-00736920/document

18. Fall, K. (2003), "A delay-tolerant network architecture for challenged internets," in *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications* (SIGCOMM '03). ACM, New York, NY, USA, 27-34. DOI=http://dx.doi.org.libproxy.mit.edu/10.1145/863955.863960.

19. Faratin, P., D. Clark, S. Bauer, W. Lehr, P. Gilmore, and A. Berger (2008),"The Growing Complexity of Internet Interconnection," *Communications & Strategies*, No. 72, 4th Quarter 2008, pages 1-21, available at https://ssrn.com/abstract=1374285

20. Harris, R.G. (1998), "The internet as a GPT: factor market implications," in E. Helpman (ed) *General purpose technologies and economic growth*, MIT Press: Cambridge, MA, 1998

21. Hess, C. (2000), "Is there anything new under the sun?: a discussion and survey of studies on new commons and the Internet," in Constituting the Commons–the eighth biennial conference of the International Association for the Study of Common Property, Bloomington, Indiana

22. Hess, C. & E. Ostrom (2007), Understanding knowledge as a commons, The MIT Press: Cambridge, MA, 2007

23. Hofmokl, J., 2010. The Internet commons: towards an eclectic theoretical framework. *International Journal of the Commons*, *4*(1).

24. Howell, C. and D. West (2016), "The Internet as a human right," Brookings Institution, November 11, 2016, available at https://www.brookings.edu/blog/techtank/2016/11/07/the-internet-as-a-human-right/.

25. Huston, G. (2017), "The death of transit and beyond," available at http://www.potaroo.net/presentations/2017-10-23-death-of-transit.pdf

26. Labovitz, C., Iekel-Johnson, S., McPherson, D., Oberheide, J. and Jahanian, F. (2010) "Internet inter-domain traffic," in ACM SIGCOMM Computer Communication Review (Vol. 40, No. 4, pp. 75-86). ACM, available at https://jon.oberheide.org/files/sigcomm10-interdomain.pdf

27. Lambino, A. (2010), "Impending tragedy of the Digital Commons?". World Bank Blog, October 25, 2010, available at https://blogs.worldbank.org/publicsphere/print/impending-tragedy-digital-commons

28. Lehr, W., D. Clark, and S. Bauer (2013), "Measuring Performance when Broadband is the New PSTN," Journal of Information Policy, Vol. 3 (2013), pg. 411-441 (available at: http://jip.vmhost.psu.edu/ojs/index.php/jip/article/view/94)

29. Lehr, W. and D. Sicker (2018), "Communications Act 2021," forthcoming in *Journal of High Technology Law*, 2018, available at https://ssrn.com/abstract=2944257.

30. Lehr, W. and D. Sicker (2017), "Would you like your Internet with or without Video," Journal of Law, Technology & Policy, vol 2017 (issue 1 Spring), available at http://illinoisjltp.com/journal/wp-content/uploads/2017/05/Lehr.pdf

31. Richter, P., F. Wohlfart, N. Vallina-Rodriguez, M. Allman, R. Bush, A. Feldmann, C. Kreibich, N. Weaver, and V. Paxson (2016) "A Multi-Perspective Analysis of Carrier-Grade NAT Deployment" in Proceedings of ACM Internet Measurement Conference (IMC).

32. Richter, P., M. Allman, R. Bush, and V. Paxson (2015), "A Primer on IPv4 Scarcity" in ACM Computer Communications Review (CCR) 45(2).

33. Sowell, J. (2015), "Finding Order in a Contentious Internet," PhD Thesis, Massachusetts Institute of Technology, available at https://dspace.mit.edu/handle/1721.1/97324

34. Stocker, V., G. Smaragdakis, W. Lehr, and S. Bauer (2017), "The Growing Complexity of Content Delivery Networks: Challenges and Implications for the Internet Ecosystem,"

Telecommunications Policy, Vol 41 (10) 1003-1016, 2017, https://doi.org/10.1016/j.telpol.017.02.004.

35. Yoo, C. S. (2016). Modularity Theory and Internet Regulation. *University of Illinois Law Review, 2016*(1), 1-62.

36. Zhang, L., A. Afanasyev, J. Burke, V. Jacobson, P. Crowley, C. Papadopoulos, L. Wang, and B. Zhang (2014) "Named Data Networking," *ACM SIGCOMM Computer Communication Review*, *44*(3), pp.66-73.

37. Zittrain, J. (2006), "The Generative Internet," 119 Harvard Law Review 1974 (2006), available at https://dash.harvard.edu/bitstream/handle/1/9385626/Zittrain_Generative%20Internet.pdf?sequence=1).