

The Political Nature of TCP/IP

Honors Thesis for Science, Technology, and Society

December 2011

University of Pennsylvania

Rebekah Larsen

Advisor: Dr. Ann N. Greene

ACKNOWLEDGEMENTS..... 3

ABSTRACT..... 3

PROLOGUE.....	4
INTRODUCTION.....	5
TCP/IP: A SHORT TECHNICAL PRIMER	8
METHODOLOGY.....	16
HISTORY: THE EVOLUTION OF TCP/IP AND ITS ASSOCIATED POLITICS.....	18
BEGINNING BACKGROUND (1960s).....	21
INVENTION/FORMATION (1973 - 1983).....	26
COMPETITION, POPULARIZATION, AND POLITICS (1983 - 1992).....	30
TCP/IP DURING THE INTERNET EXPLOSION AND COMMERCIALIZATION (1992 - 2000s).....	36
THE MATURATION OF A SYSTEM (TCP/IP TODAY).....	40
CONCLUSION	47
GLOSSARY.....	48
BIBLIOGRAPHY.....	49

ACKNOWLEDGEMENTS

Since this project was inherently disciplinary and I am more a student of social science than computer science, I must thank the generous and knowledgeable people who took time out of their busy schedules to explain concepts, put me on never-ending new paths, and provide general encouragement. I am greatly in their debt, and extremely grateful.

Within the STSC department, I have to cite Dr. Nathan Ensmenger, who first helped me come up with the idea of studying the formation of TCP/IP. His support and careful comprehension of the jumbled questions I had helped me shape this project into something tenable. I am also grateful for such a wonderful advisor, Dr. Ann N. Greene, whose endless responsibilities in the department don't detract from the quality of her advice and feedback.

Within the wider umbrella of the University of Pennsylvania, several other individuals made themselves available for general questions, discussion, and invaluable insight. Penn Law professor Christopher Yoo provided wonderful sources, inspiration, and an essential economics perspective, professor Kevin Werbach of the Wharton Legal Studies department provided vital contacts to industry leaders and new venues of information, professor Peter DeCherney of the Cinema Studies department provided contacts within the academic world, and Engineering professor Roch Guerin provided historical and technical background.

Finally, I would like to thank the Internet Society for hosting a very informative INET Philly conference, and particularly two leaders from Comcast—Richard Woundy and Jason Livingood. Their industry expertise and experience with the historical and contemporary formation of TCP/IP (and willingness to patiently describe much of it to a non-engineer) was also indispensable.

ABSTRACT

Despite the importance of the Internet in the modern world, many users and even policy makers don't have a necessary historical or technical grasp of the technology behind it. In the spirit of addressing this issue, this thesis attempts to shed light on the historical, political, and technical context of TCP/IP. TCP/IP is the Internet Protocol Suite, a primary piece of Internet architecture with a well-documented history. After a technical overview, detailing the main function of TCP/IP, I examine aspects of the social and developmental record of this technology using STS theoretical approaches such as Hughesian systems theory, Social Construction of Technology (SCOT), and Langdon Winner's brand of technological determinism. Key points in TCP/IP evolution, when viewed from an STS perspective, illuminate the varied reasons behind decisions and development of the technology. For example, as detailed in this paper, both technical and political motivations were behind the architectural politics built into TCP/IP in the 1970s, and similar motivations spurred the rejection of OSI protocols by Internet developers two decades later. Armed with resultant contextual understanding of previous TCP/IP developments, a few possible directions (both political and technical) in contemporary and future Internet development are then explored, such as the slow migration to IPv6 and the meaning of network neutrality.

PROLOGUE

This semester's project began as a rather aimless search for answers to some of my much larger questions about technical expertise and politics, especially when it comes to Internet and computer technologies. Eventually, as large and ambitious an aim as it sounds, I hope to someday have some glimmer of an idea of *how* literal code becomes law—I want more understanding of how the Internet shapes (and is shaped by) its inhabitants, who are tailored to a physical world with all of its constraints and more recognizable barriers.

I decided to study the bones of the Internet, TCP/IP, with the mindset described in the oft-repeated words of computer scientist Mitch Kapor: “Architecture is politics” (Kapor 2006). It's easier to understand exactly how powerful this idea is when applying it to a space such as the Internet—an architecture born out of a military context and being used by billions in innumerable ways, yet possibly deeply shaped by just a few at its beginning and at crucial times (Miniwatts Marketing Group 2011). I wanted to know more about technical experts' direct and indirect influences and eventually, how their personal political views (if politics refers to organizations of power and the way collective decisions are made) manifest in their work. But I needed to start at the beginning. This paper is the result my trying to familiarize myself with a complicated technology, both in technical, social, and organizational terms. This historical and general research is presented in such a way to support the somewhat broad thesis that protocols are political, and in more ways than one.

INTRODUCTION

Few people today, especially those living in industrialized societies, have not heard of the Internet—and most likely, just about every human being on the planet has in some way been affected by this giant interconnection of computer networks. The Internet is a booming place of global commerce, a portal to worlds of entertainment, art, communities, and knowledge. It has become such an integral part of life for billions around the globe that quotes such as this one (attributed to computer scientist Guy Almes across the web) isn't so far off the mark: “There

are three kinds of death in this world. There's heart death, there's brain death, and there's being off the network.”¹

For the generation born into such a connected world, it would be hard to imagine this network as anything other than the omnipresent, democratic connector it is seen as now. But in the early days of the Internet, there was a switch controlled by network builders that could turn the network on and off. Vint Cerf, a veritable father of the Internet, was one of these builders, and he later recounted using this switch to force users to upgrade key parts of their networks—by turning off the outdated parts of the network for days at a time, he and other network controllers were able to force users into making the demanded change (Odom).

The idea of such centralized control and the sheer absurdity of turning off the Net seems almost ludicrous today. Why is that? As a large economic and technological system, the Internet encounters problems that would welcome centralized control. One of the reasons that such an idea seems out of place (especially for users in nations that emphasize openness and democracy) is the type of politics that are often associated with the Internet, as well as the importance of the network that should not be turned off for a minute, let alone a day. This paper is an attempt to explore the evolution of some of these politics by examining the development of the Internet's inner intelligence: its protocols.

The Internet could be described as political on more than one level. The normal user understands this from usage and content; this big abstract web of connections carries around rhetoric and content from every view and every soapbox possible. One can take a step further and look at the control of the Internet as a political tool; current debates such as net neutrality (which will be addressed later) are essentially line dividers, leading to arguments for freedom in two different senses of the word—maybe freedom of the market or freedom in terms of government-mandated equal access. But what if we go a step further and farther back in time, and look at the

¹ Despite an hour searching online for the official source of this quote, I was unable to find anything except a lot of repetitions of the attribution. This is one example of one problem with the enormity of the Internet—the possible decontextualized nature of information. Here is one typical online source that quoted but did not reference the quote: <<http://www.quote garden.com/computers.html> >

construction of the Internet? Can one find politics in the construction and adoption choices of various groups, and then see those politics in the actual workings of the Internet today?

Fortunately, there is a way to trace politics in network development: through the technical standards that arise and compete with each other. After all, “in most cases, standardization is a means by which various social groups realize their interests” (Kim et. al, 282). Computer protocols are like any other kind of protocol—they are standardizations of certain actions, meant to unify actors and create a functional understanding of a process. One can think of the standards for mailing a letter (sender’s formatted information in the top left corner, receiver’s in the center, stamp in the top right corner) or other communication protocols in various cultures (kisses on the cheek, perhaps a firm handshake, a bow) as examples. These protocols have grown from contexts that can be political. For example, protocols can maintain power by sheer demonstration—the ancient Chinese etiquette of mandatory kowtowing before the emperor for an audience is one such example.² But technical standards, which make the Internet possible, also have great political potential because of the way they shape environments like the Internet. As Yale scholar Laura Denardis puts it, “[*Protocols*] control the global flow of information and make decisions that influence access to knowledge, civil liberties online, innovation policy, national economic competitiveness, national security, and which technology companies will succeed” (6).

In other words, control of protocols can lead to a deep influence within the listed areas provided by Denardis. By that same token, protocols (and the arguments and people behind them) can be used to gain insight into current power structures. They are, after all, still established by humans—and there must be a consensus for a protocol to function. Though these are postulations much too broad and lost in the infinite realities of history to specifically define (i.e. exactly how and to what degree protocols have shaped users and vice versa), it is to be hoped that this paper

²² Kowtowing is a series of standardized bows and kneels required before ancient Chinese emperors. Proper kowtowing was recognition of the emperor’s (and China’s) superiority, and many South East Asian countries observed such protocols to trade with the empire. This is just one example of how protocols are political standards, even in communication between humans (Urban 2011).

will shed some light on how protocols can reveal politics as well as maintain and embody them.

Thus, the rest of this paper is organized as a technical then historical overview of TCP/IP, following a thread of examples that hint at and even illuminate politics in the TCP/IP system. There are many actors, architectures, and organizations; I have compiled a brief list of acronyms for reference at the end. Understanding the basis of the internetwork we've come to recognize as cyberspace requires some in-depth background, but to perceive the politics one must know the architecture.

TCP/IP: A SHORT TECHNICAL PRIMER

In 1973, Vint Cerf wrote to other engineers working on TCP, describing its main function as the following: *“The job of the TCP is merely to take a stream of messages produced by one HOST and reproduce the stream at a foreign receiving HOST without change”* (McKenzie 2011). As one of the long-time inventors of Transmission Control Protocol/Internet Protocol (referred to as simply TCP in the beginning), Vint Cerf wrote a description that holds true today. That's one of the most notable things about the basic Internet Protocol Suite—it's been robust enough to support enormous changes in scale, as Internet usage has gone from hundreds to billions of users. From its first major uses in the early '80s until now, TCP/IP has been helping to move and assemble data around the world in a remarkably reliable manner (Kozierok 2005).

How does TCP/IP coordinate the movement of this information? Between networks, or groups of computers that cluster together, talking to each other in the same way and using similar hardware and languages. In other words, TCP/IP facilitates communication on the Internet, the giant connection of millions of networks. TCP/IP does not use circuit switching, which was the primary way of moving information on telecommunication networks (e.g. telegraph and telephone networks) prior to the 1960s (Abbate 2000, 8). Circuit switching is at its heart an opening of a direct pathway between two hosts. No matter how it is established, the network maintains this line between the two hosts involved in communication for the duration of the 'conversation'—just as when talking on the telephone (Kozierok 2005). Yet this kind of communication has been seen as

inefficient as well as vulnerable, as will be described later in this paper. One solution was packet switching. In a network operating on this principle, the data is packaged into separate pieces and sent out over a non-specified path. (See Figure 1.) In other words, messages get chopped up, and each portion travels on its own with a little address tag attached to it until it reaches its destination.

TCP/IP operates using packet switching. In a very fundamental sense, TCP/IP is just a bunch of rules that specify exactly how communication can happen between computers by breaking up these messages and moving them around. As explained by Douglas Comer, “In a sense, protocols are to communication what algorithms are to computation...a communication protocol allows one to specify or understand data communications without depending on detailed knowledge of a particular vendor’s network hardware” (Comer 1995). In other words, protocols in this case are also layers of abstraction: they allow users or applications to employ big, powerful tools without dealing with the little bits that compose those tools—almost like a driving a car or operating a buzz saw without manually moving every component at a time.




Figure 1: As packets move along nodes of the network, they have a variety of pathways they can take to reach the destination. “Computer History in a Timeline.” *Akasha*. 20 April 2011. <<http://www.goodfortheearthideas.com/Akasha/computerHistory/ctime19601980.html>>, Accessed March 2011.

TCP/IP itself is a layered set of protocols, each of them dealing with a lower (or higher as traversed) level of abstraction. Before I go any further, I should clarify that TCP/IP are the names of the two most important (expansively connective) protocols in what is generally known as the Internet Protocol Suite—this is what makes up the protocols within what we know today as the Internet. Funnily enough, the IPS is normally referred to as TCP/IP, so I will describe the general

four protocol layers that make up the IPS but I will go into a bit of functional detail when it comes to the actual Transmission Control Protocol and Internet Protocol.

The five layers of abstraction in TCP/IP, from the most abstract layer to the least, are generally known as the application layer, the transport layer, the Internet (or network) layer, the data link layer, and finally the physical (or access, for the purposes of this paper) layer (Yoo 2011, 11-15) (Braden 1989). Each layer (except the Internet layer) can and does contain a number of protocols that perform certain functions designated to that layer, but in different ways. The Application layer protocols are those that allow communication between computers when performing what we see as traditional tasks, including SMTP (Simple Mail Transfer Protocol) for email, FTP (File Transfer Protocol) for file transfer and Telnet for remote login (Hedrick 1987, 1). These, among all applications protocols for moving particular kinds of information on layered networks, all rely on the lower level functions provided by the other layers. They essentially take for granted the transmission capabilities the next level provides, handing off whatever information they wish to move in a designated format to TCP.




Figure 2: This illustration shows the conventional four-layered Internet Protocol Suite, and the direction that layers communicate when moving data from host to host. “How Encapsulation Works Within the TCP/IP Model.” 27 Jan. 2008 < <http://learn-networking.com/tcp-ip/how-encapsulation-works-within-the-tcpip-model> >, Accessed March 2011.

The Transport layer contains middle level protocols in the Internet Protocol Suite, such as

TCP. As put by Rutgers researcher Charles Hedrick, TCP “is responsible for breaking up the message into datagram’ s [sic] (those pieces used for packet-switching), reassembling them at the other end, resending anything that gets lost, and putting things back in the right order” (5). TCP is a set of rules used to manipulate these datagrams before and after they travel out along the network’ s lines using IP. Datagrams are commonly referred to as packets, though the two don’ t really mean the same thing—a packet is “physical” , appearing on wire, while a datagram just refers to a unit of data (4). But in this paper, the terms are interchangeable.

‘Internet’ is the next layer down, where the IP resides as the sole protocol. IP take the packets that TCP (or other protocols on this layer) hands it and transmits them according to the address that TCP stamped on the packet. IP must function with each host having its own unique address in a global setting (Comer 1995, 59). In a nutshell, the job of IP sounds simple but when it works with large, complex networks, the IP (in conjunction with routing tables) takes care of multiple routes and compatibilities (Hedrick 7).³

The next layer of the Internet Protocol Suite, the data link layer (also including the physical layer depending on the representation) contains protocols that determine how packets (known as frames at this level) move about in a particular autonomous system.⁴ The Internet (even by name) is a great conglomerate of networks with different sizes, protocols, and even different materials. For example, one of the most popular networking technologies today, Ethernet, is a broadcast protocol on this layer (Yoo 2011, 14).

The final layer (physical or access) is where much of the abstraction is stripped off and even more of the uniqueness of the specific network comes into play. This layer provides an interface between the IP’ s commands and the actual hardware doing the movement (Denardis 2009, 7). These are the protocols that allow transmission of datagrams on a particular network.

³³ “A routing table is used by TCP/IP network routers to calculate the destinations of messages it is responsible for forwarding. The table is a small in-memory database managed by the router's built-in hardware and software.” Depending on the size of the network, the computers that store these routing tables must keep track of just a few to hundreds of thousands of possible routes, calculated with a list of IP addresses (Mitchell 2012).

⁴⁴ Autonomous systems (ASes) are network organizations/companies that maintain collections of routers. They can be large backbone providers, consumer-facing providers (such as Comcast), and organizations such as Penn that provide access to members. As entities with different goals, there are economic relationships based on traffic flow between these systems—for example, Penn pays an regional network conglomerate named MAGPI for Internet access, and MAGPI probably pays a backbone network to send traffic over long distances (Kearns 2011).

There are a variety of technologies that are possible on this layer, seeing as how there are a variety of ways to move data and maintain networks. But since this is beyond the scope of this paper, which focuses on the layers that make movement between the networks possible, any questions are deferred to Chapter 2 of the popular engineering text *Interconnections* by software engineer Radia Perlman.

As stated before, the two layers of most importance in this paper are the Transport layer and the Internet layer, and TCP/IP are the protocols (for the most part) that connect the multiple networks that make up today's Internet. TCP is involved in prepping and controlling packets and IP provides the means to travel. I will walk through their functions in a little more detail, using simple text graphics slightly modified from those used by researcher Charles Hedrick in a 1988 summary of TCP/IP for students. His illustrations demonstrate how the packets are segmented, sent, received, and checked through protocol, and the basic functions of TCP/IP he explains are much the same today as they were in 1988.

First, imagine a data stream, some information a computer wants to send to another, as this: *****. The TCP will first chop the data into pieces, as large as possible after considering the capabilities of the networks that are involved. So now the stream looks something like this: *** *** *** *** *** **. Next, the TCP will put its own little message at the head of each of these pieces: T*** T*** T*** T*** T*** T***. This header contains the information that allows TCP to reassemble the packets in order at the receiving end, as well as check for missing or damaged packets. Packets in this format are sent back and forth between the communicating machines, and through them the machines tell each other if messages were received, the correct ordering of packets, and the parts of the actual communication. Of course, there are a lot of nitty-gritty details when it comes to this process (such as the other pieces of the header, the port specifications for making virtual circuits, and some of the special functions for speed of transfer), but in a nutshell, this is what TCP does (Comer 1995, 192 - 198).

Imagine these packets at the time of deployment from TCP: T*** T*** T*** T*** T*** T***. They are then handed off to the Internet Protocol layer, and the IP will only concern itself

with the destination address in the TCP header, and so after taking a look, it adds its own header: IT*** IT*** IT*** IT*** IT*** IT***. This header has bits of information like the source and destination addresses, the type of protocol that handed the packet off (remember, IP can support various other protocols as a lower-level protocol), and more information about the packet that allows IP to make sure that nothing is lost or damaged in travel. For example, there are other parts of the header such as the flags and fragment offset that keep track of packets' contents if they have to be subdivided again (Hedrick 1987, 6 - 7). Yet the simple, unassuming nature of IP as a carrier and nothing more is a good description of its overall capabilities. So together, TCP/IP is what allows information to be transferred in such an all-encompassing manner, geography and differing network standards aside.




Figure 3: An illustration of the various headers that are tacked on at each layer in the IPS.
 "How Encapsulation Works Within the TCP/IP Model." Learn Networking, 27 January 2008.
<http://learn-networking.com/tcp-ip/how-encapsulation-works-within-the-tcpip-model>




Figure 4: An illustration of the five layers of the Internet Protocol Suite, and the direction that layers communicate when moving data from host to host. The routers in the middle don't use the application or transport protocols—they only deliver the message using IP on the network layer. Colorado State University, CS 551: "Computer and Network Protocols." 2001. Accessed October 2011.
< <http://www.cs.colostate.edu/~cs551/CourseNotes/Communication/NetworkProts.html> >

When each of the protocol layers completes its tasks, the message (and what's left of the added header) is handed up to the next layer, until the message is assembled and fully communicated. But where does all of the handing back and forth take place? Figure 4, in addition to the following paragraph, illustrate how the middle of the network is a simple transporter, running just the lower layers of the Internet Protocol Suite, while the ends of the network experience more responsibility and more diversity.

For the purpose of this paper, hosts could be called basically anything that is hooked up to the network and has its own IP address. Hosts gets a little more complicated than that—for example, there are computers maintaining networks that are called IP hosts, which manage addresses for distribution among possibly many computers on their network (Padlipsky 1982). But in this context, one could think of hosts as computers that are requesting/receiving information from another host, and the routers are the middlemen, the computers in the center of the system that convey messages. When it comes to the layered protocols, the routers run the bottom-layer protocols—they do not deal with the transport layer or the application layer. We can think of them as simply moving the packets, no questions asked, between the smarter hosts, who deal with reassembly, error control, formulating a response, breaking it down, and sending it back out on the simple singularly functioning middle.

The resulting distribution of responsibilities is often described as a 'dumb network', and the implications are illustrated in this hourglass shape. The IP, as the uncomplicated middle created by the routers' topmost layers, is kept deliberately simple because it allows the ends of the network to diversify (Zittrain 2008, 67-69). In other words, because the IP keeps its demands and

function to a minimum, a number of applications as well as network technologies can use this transfer system to communicate. TCP (and protocols like UDP, User Datagram Protocol) operate on the layer that provides network command and control, taking the service out of the middle. The hosts, running these upper layer protocols, have to reassemble packets as well as do damage control—determining as best they can from the fringes what went wrong if a packet doesn't arrive, and resend or adjust their queries for congestion.




Figure 5: An abstract visualization of the innovation/diversity that is allowed at the ends of the network because of the simplicity of its middle. Taken from Jonathan Zittrain's *The Future of the Internet*. Searls, Doc. "Understanding Infrastructure." *Linux Journal*, 19 April 2008. <<http://www.linuxjournal.com/content/understanding-infrastructure>>

This is a deliberate design decision; there have been other networks (as will be illustrated in the historical overview) that were designed to provide more control from the middle of the network, as some postulate, at the cost of innovation on the ends (Wilson). As it turns out, the centralized control of the network has also been a hot political topic for decades, reappearing in various

historical stand-offs and clashes between builders/users of the network. In fact, the birth of TCP/IP was the result of a power change.

But how do we elucidate the politics behind and built into this system? I've decided to interrogate history with some STS theories; they might not definitively outline the story and politics behind TCP/IP, but they shed light on certain areas of development, and ultimately show that protocols and their formation are deeply political in more ways than one.

METHODOLOGY

The reason I've tried to detail some of the functionality of TCP/IP is of an interrogative nature: in order to understand an invisible architecture's influence and possible politics, one must have a functional understanding it. Penn Law professor Christopher Yoo has alluded to as much in a presentation entitled "Layering, Modularity Theory, and Internet Policy", detailing how current Internet policy directives from non-engineers relies on less-than-working knowledge of Internet functionality—the result is vague, hard-to-implement policies that many engineers no doubt shrug their shoulders at.

Even Eric Schmidt, the head of superpower search engine Google, has recently stated how little many members of Congress understand how the Internet functions, and then try to implement what seems like impossible directives. "For every one of these Internet-savvy senators, there's another senator who doesn't get it at all. And it's not a partisan issue. It's true in both parties" (Cunningham). And this lack of understanding can create unrealistic expectations. As Schmidt put it, regulation can "prohibit real innovation, because regulation essentially defines a path to follow—which by definition has a bias to the current outcome". Internet engineering, according to Schmidt and many other engineers, is a fast and dynamic process—it can't be directed in the same way as other social policies. Yet the interaction between technology and outer society (as well as the inner functions of both) must be recognized and understood as much as possible for any sort of effective policy to occur.

Trying to get some idea of the mechanisms behind technology and society interaction is a daunting process. There are a variety of STS theories available as tools for interrogation, and I

though I do not intend on focusing on one theory, I will attempt to show the variety of ways politics are a part of protocol systems through several approaches—whether through the arrangement of the technology itself or through the politics surrounding its governance and development. This is meant to show the variety of ways that technologies and politics *can* interact. There are three main theories I will allude to throughout the upcoming section of TCP/IP history: Social Construction of Technology (SCOT), Langdon Winner's artifact politics, and Thomas Hughes' systems theory.

SCOT in the history and sociology of technology was popularized in the 1980s, in part by Trevor J. Pinch and Wiebe E. Bijker, both sociologists in science and technology (Bijker 1987, 1). In their 1987 joint essay on social constructivism, “The Social Construction of Facts and Artifacts”, they detailed a “multidirectional” model for technology selection and development, showing how artifacts are contingent on social factors, and how the “successful stages in development are not the only possible ones” (28). Thus, SCOT dismisses the notion of fixed linear development of technology, and instead posits stages of development and research as well: interpretive flexibility, closure, and ‘relating the content of a technological artifact to the wider sociopolitical milieu’ (40). These will be explored later in this paper, with TCP/IP as the technology in question.

On the opposite end of the theoretical spectrum is technological determinism, which points to the possibility of technology shaping society. In 1986, Langdon Winner published an influential article read by every subsequent history of science and technology student. “Do Artifacts Have Politics?” asked Winner, and according to him the answer was yes. He claimed, “systems of modern material culture can embody specific forms of power and authority” in two major ways: through arrangement/adoption or through inherent compatibilities with certain political relationships. For the first, he gives a compelling example of the discriminatory politics constructed into the very low overpasses of New York's Long Island; those lower on the socioeconomic ladder (many of them non-whites) could not afford travel to the beaches except by bus, and the bridges presented a very solid, architectural denial. The other kind of system-embedded politics Winner described is

more controversial because of its deterministic nature; he argued that some technologies are strongly compatible with and even demand a certain “set of social conditions” (Winner 1986, 130). One powerful example is the nuclear power—if a society is going to maintain this kind of system, its very dangerousness requires a “centralized, rigidly hierarchical chain of command” over the technology—something with which democratic societies might struggle (131).

Thomas Hughes’ systems theory is an attempt to reconcile these opposing ideas of social vs. technological determinism by incorporating both at different times during technological development. In “The Evolution of Large Technical Systems” (1987), Hughes details how technology is always embedded a system, and a young system is deeply shaped by its surroundings, often “inaugurated” by “radical inventions” (Bijker 1987, 62). Hughes argues that these systems (which are often embedded in larger and larger systems as well) follow a loose pattern that changes the direction of influence in the system as it matures (56-57). In other words, technical systems tend to be socially constructed in the beginning, but after rough stages of development, competition, and consolidation (among others), these systems tend to develop a kind of technological momentum (57-59).

This is Hughes’ incorporation of technological determinism—these systems “have a mass of technical and organizational components; they possess direction, or goals; and they display a rate of growth suggesting velocity” that look something like autonomy to observers (76). Hughes also details how ‘reverse salients’, road blocks arising within a system itself as certain components advance and outdate other components, can disrupt this pattern of development. If they cannot be overcome by the system, a new system via radical invention can be born (73-75).

I employ these three methods of inquiry at choice times in the following historical sketch because each provides a distinctly different way to view how choices are made about technology. Some choices (and their binding effects in the technology) have already been alluded to in the previous sections (e.g. the centralization or distribution of network control), but this next section aims to shed light on the context of these choices, using insight from Hughes, Pinch, and Winner.

HISTORY: THE EVOLUTION OF TCP/IP AND ITS ASSOCIATED POLITICS

Much about TCP/IP has remained remarkably the same in the past 30 years (Huston 3). Yet this doesn't mean that anything about the Internet is static—on the contrary, the relatively stable existence of TCP/IP is something of an anomaly when compared to other Internet components, and has both engineers and policy makers continue to argue in big forums and to big audiences whether this is a good or bad thing.⁵ The following quotes, though neither giving a comprehensive picture of Internet development culture or projection, still illustrate some relatively constant themes in TCP/IP—and overall Internet—evolution.

Speaking from a telecommunications point of view, network consultant Thomas Noelle articulated one side of a recurring tension that pops up over and over in recent networking history.

“The Internet is an important cultural phenomenon, but that doesn't excuse its failure to comply with basic economic laws. The problem is that it was devised by a bunch of hippie anarchists” (Nolle 2001). Netheads vs. Bellheads, TCP/IP vs. OSI, X.25 vs. vendor networks—all were contexts for disputes about control and management of networks, and will be addressed in later portions of this paper. These disputes were also deeply tied to the technical standards of the Internet, and the way information is routed. This tension, which is about power at its heart, is deeply political in nature, and has led to some interesting characterizations of technologies and groups involved in the TCP/IP evolution. This quote also touches on some of the perceived embedded principles of TCP/IP, which could be examined using Winner's ideas of arranged politics.

Speaking from the 'hippie' side, computer scientist David Clark coined this mantra of the Internet builders in 1992: “We reject kings, presidents, and voting. We believe in rough consensus and running code” (Hoffman 2011). It has characterized much of the subsequent discussion that dominates the Internet history, painting it with ideas of utopian development. But

⁵ Columbia Law professor Tim Wu coined the term 'net neutrality' in 2003, describing one side of current debates about the way information flow should be controlled on the Internet. Net neutrality advocates generally argue that the government mandate ISPs (internet service providers) to treat all every kind of packet equally, not discriminating based on the user or type of transmission—thereby supposedly preserving a platform for innovation. Those against such measures cite lack of market freedom and the possible backlash to consumers if flexible pricing is banned (New York Times 2010).

was it really all about the perfect environment? There's little doubt the Internet origins have gained this reputation, such as how it was ultimate meritocracy from the beginning, where nerds were born and thrived, where "no one knew if you were a dog"⁶ and if a user gave a suggestion, it would only be taken for its intrinsic value and not its source. But like any other technology, TCP/IP was constructed and maintained with political aims, and itself embodied politics that weren't necessarily egalitarian. After all, looking closely at the David Clark quote, it's evident that engineers rejected that mechanism of democracy and voting, probably for the "partisan compromises" and bureaucratic structure they imply (Russell 52 - 53).

Law professor Christopher Yoo pointed out in February 2011 that the Internet arose from an experimental context: "[The Internet's] layered model evolved during the design process through experimentation and compromise rather than a precommitment [sic] to a particular set of principles. [This process] of implementation and experimentation is responsible for much of its success" (Yoo 2011, 25-26). Its original intended purposes were nothing like what it is used for now, and perhaps this malleable nature still exists. Could this also mean that Internet governance is something not set in stone as well? Such an analysis follows from the main thrust behind SCOT: there are multidirectional paths technologies can take, dependent on their incumbent and shaping social groups.

Echoing the sentiments of Christopher Yoo but in a wholly technical setting, Internet pioneer Jon Postel wrote this introduction to a popular network engineering text in the 1990s: "Like all communications systems, the TCP/IP protocol suite is an unfinished system. It is evolving to meet changing requirements and new opportunities" (Comer 1995, xxi). His description of TCP/IP imparts to the technology a sense of autonomy, evolving as a system in response to contextual changes. Such rhetoric can indicate that the technology seems self-determining even to its original creators—similar to Hughes' description of the momentum of established systems. TCP/IP has not remained wholly the same over the years. Perhaps its basic

⁶ In the July 5, 1993 issue of *The New Yorker*, Peter Steiner published a cartoon of two dogs discussing the beautiful anonymity the Internet can provide. The result was a catch phrase for discussion of the various implications of online anonymity (Steiner 1993).

capabilities have stayed consistent, but the uses of the Internet, the people who are involved with it, and the organizations behind its design have been in flux. This next section will attempt to examine both major technical and contextual evolutions of TCP/IP from its roots to the present day.

The history and projection of TCP/IP here will roughly be divided into five sections: the setting and precursors, the invention and formation of basic TCP/IP, its competition and popularization, TCP/IP during the commercialization of the Internet, and a look at some current political and technical developments around TCP/IP. Each section, though overlapping in terms of chronology, will attempt to historically examine some of the key players shaping TCP/IP, and the way TCP/IP has shaped them. (Goldstein 278)

BEGINNING BACKGROUND (1960s)

Paul Baran, as a young engineer, went to work for the Rand (Research AND Development) Corporation, a non-profit think tank funded by the United States government in the 1950s. Attracting the best minds with autonomy, high salaries, and the ability to influence policy, Rand was a go-to for many technology projects (Hounshell). The Cold War was gearing up; this was the time of ‘hair-trigger nuclear ballistic systems’. And Baran saw a solution to the need for survivable communications—a ‘distributed system’, or what soon became known as packet switching (Abbate 11 - 13).

Two men in particular saw this solution around the same time, and for different reasons.⁷ Baran was the first, and decided to devote time to the concept of a network that “was independent of centralized command and control” because of worries about a nuclear attack (Galloway 5). Packet switching, as a technology that allowed pieces of messages to independently find their way to their destinations, was a key feature in his design of 1964. Message switching was nothing new, but it had never been proposed in such a way before—Baran wanted intelligent

⁷ The history of science is peppered with multiple simultaneous discoveries, leading to the idea societal understanding reaches a point where the next breakthrough is primed for discovery. Famous examples include the simultaneous discovery of calculus by Sir Isaac Newton and Gottfried Wilhelm Leibniz, the theory of evolution of species by Charles Darwin and Alfred Russel Wallace, and the multiple discovery of oxygen by more than three people in the 18th century. The age of computers has made this phenomenon even more pronounced in the history of technology—such as with the simultaneous invention of packet switching (Goldstein 278).

nodes, computers rather than telephone switches, that could store the packets and forward them at appropriate times, in the right direction. There would be a cost on the system: break up and reassembly of the messages would be time-consuming and difficult (Abbate 2000, 10). The tradeoff seemed possible in such a Cold War context though—robustness, flexibility, an inability to eavesdrop, and actually more efficient use of the network as a whole all combined to make the proposals seem feasible. But the plan for Baran’s network was scrapped at his request in the 1960s. The implementation had been shuffled to the newly created Defense Communications Agency; Baran and his backers knew that the organization didn’t have the technical ability to pull off such a network, and so they pulled out to protect the reputation of the design. Baran had enough control over this idea, in a research yet military context, that he could control its physical existence (21).

Around the same time, another researcher across the pond was working on his version of packet switching in a much different context. Donald Davies was working at National Physical Laboratory in the United Kingdom. The UK was in the midst of fathering technological growth in an effort to stay abreast with the rest of the world, and Davies saw packet switching as a way to utilize more resources. Time-sharing was the big thing, and in this era of big computers and big prices, the ability to share computing power at multiple sites was extremely intriguing.⁸ Davies also thought of packet switching as ultimately a commercial product, where “packet switching networks would compete with other communication systems to attract and serve the business users” (29).

Yet the networks that Davies ultimately built could hardly have been called such. The Mark I had one node, and though its follower, the Mark II, stayed in service at NPL in 1986, it never reached the heights that the eventual United States packet switching network did. But the context for its development was again something unique. When the Defense Advanced Research Projects Agency of the United States funded the famous ARPANET, it was officially called the

⁸ In the early days of computing, most computing power was located in mainframes, huge, expensive computers that were best used by sharing. The prohibitive costs brought about renting of systems as well, and brought about ideas of sharing resources using packet-switched networks at both NPL and the ARPANET (Time Magazine 1965).

first operational packet-switched network. However, it was not built to withstand nuclear attack or provide fail-safe communications in the event of nuclear holocaust—this was a research and also a resource-sharing tool, but the legacy of Baran’s survivability goals lived on its design (Galloway 5).⁹

It was 1968 when the plans for ARPANET were submitted to ARPA by program manager Larry Roberts, and came back with a \$2 million in approval to seek out contractors. The bid was given to Bolt, Beranek, and Newman, a firm with ties to the Information Processing Techniques Office (a part of ARPA) and also to Honeywell, a computer manufacturer that ended up providing the nodes for this new network. By the end of 1969, researches had installed the first four nodes to ARPANET (Abbate 39 - 57).

ARPANET also provided a place for development of layered, packed-switching protocols—the first instantiations were rudimentary to be sure, and not robust enough to handle an expanding network, but they no doubt influenced later forms of encapsulation. The transport protocol was called NCP, and it ran on the hosts connected through IMPs (Interface Message Processors), the Honeywell DDP 516s (a type of computer) that BBN procured as nodes for the ARPANET (Yoo 2011, 9-12). NCP stands for Network Control Program, though often mislabeled as Network Control Protocol because it essentially was protocol. A host would connect with another host via a message through NCP called a handshake, and then the two would establish a connection for communication. NCP was the protocol between the hosts, overseeing the transfer of information (Yoo 2011, 15-17).

The nodes of the ARPANET were ‘smart’—they maintained the state of the network. This means essentially that if a node was lost, a whole connection could be lost (31). Recalling the earlier technical section, TCP provides communication between hosts on the edges. According to

⁹ The influence of Baran’s fail-safe communications system goals on the ARPANET has been somewhat debated—Abbate intimates that Baran’s and Davies’ ideas of packet-switching were both used (both were consulted) during the construction of ARPANET (36). But ARPANET program manager Larry Roberts has said that ARPANET and the idea of a “secure defense network had nothing to do with each other.” It is interesting to note that differing histories can be used to prop up a certain idea of the Internet. For example, Roberts’ quote was in the context of defending ARPANET as the “future of the science in the world”, and such an artifact could be seen in a more comforting light if not seen as a military product (Segaller 1999).

orthodox beliefs of encapsulation, the TCP header of each packet shouldn't be touched as it moves through routers running IP. In this way, the actual nodes in the network are blind to congestion and/or failures—the smart ends of the network take responsibility, and the 'dumb' middle can do its job simply and efficiently (Clark 1988, 108).

With NCP and the first incarnation of the ARPANET, there was no dumb middle—the IMPs as minicomputers kept copies of packets until receiving confirmation that the packets had successfully been sent to the next node on their journey to the destination host (Abbate 67). This was a network that, in contrast to those characterized by TCP/IP, placed much more emphasis on reliability and control of the network in the middle. Scientists were the ones using and creating the network in an effort to share resources, and these two reliability and control were priorities.

The year was 1970, and the Network Working Group was the name for some individuals informally outlining the NCP protocols and discussing network architecture. Loosely brought together in 1968 by various DARPA project leaders such as Vint Cerf, the group began to guide the process of creating the ARPANET. They built a system for developing and publishing their consensus (still in existence today, known as Requests for Comment), and the organization has served as a template for current Internet protocol setters such as the Internet Advisory Board, ISOC, and especially the Internet Engineering Task Force ("30 Years..."). Interestingly enough, this was an organization that almost didn't make it past adolescence. Early on in ARPANET development, Lawrence Roberts considered turning the NWG over to a professional research group. But even at this stage, it was apparent that the sense of community and the users-as-creators aspect of networking were beneficial. He decided to let the users work their way through their network (Abbate 2000, 73).

The technical details of the ARPANET illustrate the 'rough consensus and running code' nature of this precursor to the Internet. During 1969, the NWG spread the word using RFCs, icons for the sheen of community effort that has glossed many conceptualizations of this new space. The RFCs contain many snapshots of Internet culture and development from their beginning, as well as documentation of the evolution of TCP/IP. The very name itself—*Request*

For Comment—does not denote authority or control. These are electronic archives documenting Internet technical development, from protocols to technical procedures to curious jokes inside the large community that has constituted Internet expansion since 1969 (Denardis 2009, 26).

The researchers and graduate students who first began to flesh out ARPANET (both in terms of purpose and technical specifications) wrote informal memos to each other, sending out photocopies until the memos migrated to digital format on the very network they were constructing. But today, these memos are seen as the official, formal documentation of Internet protocol standards. TCP/IP specification and tutorials reside in many RFCs, from some initial descriptions of TCP (RFC 761) to tutorials (RFC 1180) to congestion control modifications (RFC 5681). Another familiar name, Jon Postel, was the archivist of the RFCs until his death, and the thousands of documents (some very technical, some dreams of Internet future, some consisting of meeting minutes, and some just for April Fool' s Day)¹⁰ are now available online for anyone to peruse.

Stephen Crocker, another big name in Internet culture, was the first writer of RFCs and the one who tentatively entitled them as such, afraid of possibly upsetting the military funders. In a fairly recent opinion piece for the New York Times, he put into words the idea of Internet culture that the RFCs have helped perpetuate and even idealize. “This was the ultimate in openness in technical design and that culture of open process was essential in enabling the Internet to grow and evolve as spectacularly as it has” (Crocker 2009). With such descriptive rhetoric of its past as well as its history of anonymity and supposedly low barriers to entry for start-ups, the Internet has become something of a cognate for freedom and democracy in the eyes of many of its users. But TCP/IP itself has a history of power struggles embedded in its design that might make one hesitate to call the formation of the Internet a beginning of a utopia of meritocratic civil society.

But when it comes to methodological inquiry, even the first few years of ARPANET, before the advent of TCP/IP, were very influential. From a Hughesian perspective, it' s important to note that the ARPANET arose from namely military and academic systems. As Hughes puts it,

¹⁰ A techie take on the Jabberwocky entitled Arpawocky, published in 1973, was forever enshrined in Internet culture and history by becoming RFC 527. No doubt the engineer who wrote this poem had some frustrating times in his work with the early NCP/IMP architecture of the ARPANET (Merryman 1973).

“systems nestle hierarchically like a Russian Easter egg into a pattern of systems and subsystems” (Hughes 54).

The political nature of these two enveloping and overlapping systems, even if just from their virtue as cultural artifacts, would have an effect on the burgeoning ARPANET. For example, there has long been debate about the distortional effects of the Cold War and American military involvement in the direction of science. (See footnote 9.) This military distortion can of course be viewed as a push for weapons-conducive science rather than peacetime technological development, but it is often more nuanced than that. “Even in the beginning days of RAND, an organizational culture that prized intellectual curiosity and independence” began to permeate government-funded research institutions (Hounshell 1985, 242). There’s little doubt such meritocratic values were present in ARPA since Baran (as well as Davies) contributed to ARPANET development. Here the possible foundations for the meritocratic, nitty-gritty nature of the Network Working Group and subsequent organizations are visible.

One of the most powerful aspects of the Hughesian pattern for systems development is an early tendency towards social construction. Indiana University researchers Junghoon Kim and Tomoaki Watanabe wrote an excellent SCOT analysis of early Internet development, and they claim that the eventual distributed nature of the ARPANET was not a “natural course of development...nor was it a necessary requirement of the diffused network to use a packet switching protocol”. Instead, they argue that the packet switching, diffused Internet is a “product of historical contingencies” (Kim 2002, 282). The opportunities to demonstrate TCP/IP are what eventually set it apart, and those opportunities arose from the military-industrial-academic setting of ARPANET. This next historical section includes some elements that demonstrate how TCP/IP (and the Internet overall) was still affected by military aims and how consensus was created to fit certain desires for power distribution.

INVENTION/FORMATION (1973 - 1983)

The ARPANET began to grow in size and inspire similar networks quickly after it became

operational. Funding was in large supply, and researchers as users flocked to assist in development of standards and protocols (Comer 1995, 6). The first incarnation of TCP/IP came with a paper by Vint Cerf in 1973, previously quoted in the technical overview. “A Partial Specification for an International Transmission Protocol” first proposed that host computers do the work of breaking messages into chunks and reassembling them, rather than each computer (SIGCOMM 1999).

The ideas in that paper that would soon become TCP/IP were refined and published by Cerf and fellow researcher Robert Kahn, and in it they established a basic system for gateways and uniform protocols between networks (Yoo 2011, 27). The idea of networking was spreading, but for the most part, these networks were separate entities, running on their own protocols and/or medium. What kind of protocol could run on all networks, whether they used radio transmission or solid wires or even satellite broadcasting? And how would such diverse networks communicate? It was also in the early 1970s that Cerf brought together a group of researchers involved with the network to attack such questions (Abbate 2000, 127). This group ultimately came to a consensus on some important network principles—namely, that ARPANET was to become a dumb network with TCP/IP.

But this consensus was somewhat tailored. In her excellent footnotes for *Inventing the Internet*, historian Janet Abbate describes how BBN (the manufacturer of the IMPs, the middle computers of the smart ARPANET) was excluded from this first group of Internet builders, greatly in part because they would have most likely argued for a continuation of the smart network, one that provided reliability in the middle. Because of its history in constructing much of early ARPANET, other network contributors felt BBN had too much power over Internet architecture.

“Thus, a new network architecture may have been seen as a chance to re-negotiate the balance of control among ARPANET participants” (223).

Harkening back to the quotes at the beginning of this historical section, here is evidence of the experimental yet political nature of the Internet. Even before TCP/IP was built, there was a tension between those involved about how the network should be constructed—should it be smart (controlled in the middle) or open (control on the edges)? Also, here is evidence that researchers

could change major architectural principles years into its operation, though with ideas of both technical improvements and political power. It suggests a kind of consensus, one that didn't give equal voice to every view but created a majority movement through rough consensus.

These were the early days of TCP/IP, and consensus and experimentation were easier to establish because the people involved were relatively few. The malleability of the system by human actors at this point coincides with the rough pattern of technological development Hughes put forth—social construction is more evident when the system is relatively young, and these changes in the direction of the network could be spurred by political reverse salients. The exclusion of a group with differing goals also jives with a general SCOT analysis—design flexibility led to the choice of one design (and thus a certain set of goals) over the other. One can also think of this as a moment when a choice was made about the politics of a system, and then the system was arranged to embody them—one example is Robert Moses' bridges as described by Winner. But participation and technologies were growing and evolving; the system was maturing.

The year 1975 was a busy one for networking. Telenet, touted as the first packet-switched commercial network and developed by BBN, linked seven US cities by August (Mathison 1975, 24). In January, the Altair 8800 (widely recognized as the first personal computer) was introduced to the United States—and the growing ranks of hackers were gaining serious attention from the government because of subsequent security fears (Abbate 138). Three networks quickly gaining users, PRNET, ARPANET, and SATNET, were soon to become relevant parts of the history of TCP/IP. PRNET was a packet-radio network, a packet-switching network using a broadcast mechanism. SATNET was another broadcast network, but based off of satellites. And ARPANET generally used physical wires, leased from telephone companies (120-131).

Renewed interest from the military in mobile, survivable communication, specifically for the battlefield, also contributed to the push for connecting these networks (Russell 2006, 49). Networking was becoming more of a viable public technology, no longer just the area of researchers running experiments. There was a new motivation for something like TCP/IP emerging: the autonomy of networks—each could be separate and tailored for a purpose, but still connected

into a giant internetwork. In other words, the military wished to have distance from those researchers constantly experimenting on a network (Abbate 2000, 142-144). The DCA (Defense Communications Agency) took over the ARPANET project in 1975, an indication that this technology was reaching military usability and was no longer so much in the research domain. It was in 1976 that the DCA began to look for a way to connect ARPANET and AUTODIN II (a military communications network) without having each influence the inner workings of the other (McKenzie 2011, 365-367). TCP/IP (though only known as TCP at this point) was picked up. As Abbate details it, the choice of this protocol was one based in pragmatism—TCP was the only thing available (139).

Successful demonstration of TCP through those three diverse networks (SATNET, PRNET, and ARPANET) followed in 1977 (131). From here, the government launched something of a campaign and mandate for this protocol. ARPA began to fund implementations of TCP in various operating systems. UNIX was a standard operating system for university departments, and as BBN and Berkeley bundled TCP/IP capabilities and applications in its kernel, the academic world was infiltrated (Comer 1995, 7). Many understood the value of increased network membership, and early on the standards, documentation, and ultimately creation of TCP/IP were open—there were no proprietary controls on them. As put by Cerf in 1990, “TCP turned out to be the open protocol that everybody had a finger in at one time or another” (Cerf 1993). These two factors, along with the fact that the US government was providing funding for networks that conformed to TCP/IP, pushed along the TCP/IP’s growth (Comer 1995, 6). As early as 1972, the “killer application” of email was causing explosive growth on the networks running on TCP/IP (Maathius 2003).

One such reason was pressure from the United States government. Though the Internet is often historically painted as a bottom-up technology in terms of development, it’s clear that the government did have a significant effect on its direction. It was around 1982 that Vint Cerf and Dave Lynch began to strategically fiddle with the Internet switch, at the behest of the deadline set by the Department of Defense. It was a moment of SCOT “decisive closure” for the technology—some flexibility in terms of alternative protocols and alternative developers was lost at this

directive when TCP/IP became a mandate (Kim 2002, 284). Many hail January 1, 1983 as the true birthday of the Internet with the widespread movement of all government-controlled networks to the new TCP/IP standard (Jaffe 2002). Despite some military and research network separations, the Internet as a giant collection of networks held together by TCP/IP was rapidly expanding and would soon become a large part of civil society.

Throughout this time, the human organizations behind ARPANET were growing as well. The NWG began to expand and evolve, serving as the predecessor to a number of Internet architecture organizations that exist today. Cerf helped formalize the NWG into the ICCB (Internet Configuration Control Board) in 1979 at the behest of ARPA, “chaired by David Clark to assist ARPA in the planning and execution of the TCP/IP suite” (Cerf 1993). The ICCB began to expand to include researchers on the network without direct ties to ARPA. By 1984, ICCB became the IAB (Internet Advisory Board), which created a number of task forces to include the growing number of participants and keep discussions productive. The board was comprised of the heads of each of these task forces, elected by the chairman, and created a kind of “council of elders” (Russell 52).

Perhaps the best known of these task forces is the IETF (Internet Engineering Task Force). From its first meeting in 1986, the IETF has taken on a particular sheen of that ubiquitous ‘rough consensus and running code’ for its work with protocol development and its tri-annual community meetings. Over time, the IAB came to stand for Internet Architecture Board, and fell under the stewardship of ISOC (The Internet Society), a non-profit that arose in 1992 from the efforts of Vint Cerf to “provide leadership in Internet related standards, education, and policy.”¹¹ Today, there are a number of standards and architecture organizations, but ISOC and the IAB could be seen as the council of elders, maintaining large masses of network engineers producing through this process of ‘rough consensus and running code’ (Russell 52).

¹¹ As quoted from the official description on the Internet Society’s website (Internet Society 2011).

COMPETITION, POPULARIZATION, AND POLITICS (1983 – 1992)

Around 1975, when networks were barely becoming commercial ventures, various telecommunication companies (telcos) decided to create a number of public data networks (PDNs)¹². Because comparable struggles over the political power of protocols would soon reach TCP/IP and the growing Internet, a non-TCP/IP example of one of these public networks will be detailed here.

IBM (International Business Machines, an American computer manufacturer) was making much of the actual hardware for these public data networks, so they were also in a position to maintain a firm proprietary grasp on protocols. This caused the telcos and IBM to square off over standards on public data networks—the telcos wanted an open standard for obvious reasons: they didn't want to be locked in to a relationship with a manufacturer through difficulties with compatibility (Abbate 152 -157).

This is one result of proprietary protocols; they take away power from consumers by narrowing their choices in products. This dynamic led to the creation of X.25 by telcos, an interface between users and packet-switched networks that operated by opening virtual circuits (a system where the packet-switched network mimics a circuit switched environment) between hosts (Anker 2005). It gave network operators, rather than hosts, control. These public data networks were designed as public utilities. So in a strange way, X.25 embodied a goal and an opposite to the purported democratic, innovative sheen of TCP/IP: it was a smart (controllable) network but an open-standard network (Russell 52). X.25 was just one of many telcos and company-backed network constructions emerging in the '70s and '80s. IBM's SNA, Apple's Appletalk, and the virtual circuit architecture ATM (developed to transport data as well as voice) were all network technologies around the time of the ARPANET (Abbate 2000, 149).

Such concern for proprietary power was one of the reasons that TCP/IP was designed


¹² A PDN generally refers to a network that is available for use by companies and consumers outside an organization, and is usually contrasted with a voice network. Public networks in the nineteenth century were primarily data (telegraphs carrying tapped messages encoded in dots and dashes), but by the middle of the twentieth century, most of these networks carried voice. One fundamental difference between data and voice networks is the use of circuit switching for voice, and packet switching for data (When 2010, 113-115).

through ‘rough consensus and running code’ as well. The next part of this section illustrates that protocol choice, even when technical and design differences aren’t particularly different, can still be used to maintain or create certain power structures. This is a good set up for the initial history of OSI vs. TCP/IP, a story that could be more about control than about technical aspects of the protocols.

OSI stands for Open Systems Interconnection, and it is a reference model for network protocols. Essentially, it is a technique for classifying network activity into seven layers (rather than TCP/IP’s five), adding two layers (presentation and session) to the transport layer of TCP/IP. Developed through an international standards-setting organization, OSI has become the regular way for engineering students to learn about layering (Humphrys n.d.). This organization is the ISO, which stands for a word in Greek that means ‘equal’, but now officially refers to the International Organization for Standardization (Galloway 2004, 129). After being formed post-WWII to navigate the increasingly international and commercial nature of Internet standard ISO extended its influence and efforts to network protocol standards, beginning in the 1970s (Russell 2006, 52).

Many of the same tensions evident in the creation of X.25 were behind the creation of OSI. Abbate cites the jurisdictional divides between countries as well as the “complex web of interest groups” as the forces behind a “Babel of competing and incompatible standards” (152). A few examples include business network architectures such as SNA (Systems Network Architecture) from IBM, DNA (Digital Network Architecture) from Digital Equipment Corporation, and DSA (Distributed Systems Architecture) from Honeywell (Maathius 2003, 164). As the push for internetworking began in the 1970s, so did a push for openness in standards—easy compatibility

meant no proprietary standards on the shared layers. This meant that computer manufacturers would be losing market power, but consumers/users would be gaining it. OSI was just an effort to jump-start this large-scale compatibility in a top-down effort. That being said, the ISO wasn’t an organization that moved proactively—usually, it waited for a de facto standard to emerge before stamping it official (Galloway 2004, 129-130). But because of the quick development of proprietary



Not present in the model

standards (as evidenced by X.25), ISO began work on its own standards.




Figure 5: Showing the similarity in layering between OSI and the IPS. The host-to-network layer includes the data link layer and physical layer. Some TCP/IP models only display four layers. Humphrys, Mark. "OSI Stack vs. Internet (TCP/IP) Stack." Dublin City University. <<http://www.computing.dcu.ie/~humphrys/Notes/Networks/intro.html>>

OSI was first released as a simple reference model (pictured above), no actual protocols attached, in the late '70s (Russell 2006, 52). At this time, TCP/IP was already in network experiment mode, and by the time that this reference model made it through the standards committee, ARPANET and all US government supported networks were already on TCP/IP. One major difference in the rates of development lay in the opposing organizational approaches for development; TCP/IP could easily be described as a grass-roots effort, with engineers as participants in its construction rather than members. OSI came from a more top-down approach, "populated by representatives from national governments" and running on formal representational processes (Russell 52). OSI slowly developed its standards, breaking the problem into abstract-to-concrete pieces: developing a reference model, architecture, specifications, and finally actual protocols (Day 2005, 1334).

Though its slow development was a hindrance to its adoption, OSI had the might of several European nations behind it, as well as the United States (Denardis 2009, 37). The GOSIP standard (Government Open Systems Interconnection Profile), a version of OSI, was developed in 1988 for

US federal agencies; network administrators were mandated to procure compliant products by 1990. However, procure did not mean exclusively use, and like many TCP/IP network administrators throughout the US, even government employees continued to use the original Internet protocols. GOSIP was abandoned in 1994, and the action was indicative of the state of the OSI movement (Russell 2006, 53-56). As mentioned before, the OSI reference model is now best known as a popular way to teach engineering students about layering. TCP/IP went on to become the de facto standard for the Internet we know today.

“The purpose of OSI is to allow any computer anywhere in the world to communicate with any other” (Day 2005, 1335). Wasn’t this one of the key goals of OSI from the start? How did this technology, so similar that TCP could have been the official protocol for the transport layer of OSI and backed by international governmental support, not take over for TCP/IP? The answer is complex because the clash between TCP/IP and OSI has been interpreted variously.

The first and most common explanation is one of pragmatism and entrenchment. Historians James Gillies and Robert Cailliau, describing the European CERN laboratory’s involvement in the Internet, provide the mental processes of many network managers who’d committed to OSI in their book *How the Web Was Born*. The managers waited for a few years for working protocols that never came, but all of their major computers ran on TCP/IP—US manufacturers had been including these protocols to their products for years. Finally, in 1988, CERN began to run TCP/IP, becoming a major site of Internet traffic for Europe (Gillies 2000, 87-89). This was a common story with network managers, computer manufacturers, and Internet users around the world—TCP/IP was the only usable option. As put by Jon Postel, “The availability of the basic and detailed information of these protocols, and the availability of the early implementations of them, has had much to do with their current widespread use” (Comer 1995, xxvi). Such low-level user choices obviously had an effect on which technology won out. TCP/IP had contextual appeals, including its running code, but also the power of network effects was on its side. Because networks are all about communication, the more people that use them, generally the more valuable they become. Think about a fax machine—as just one machine with no connections, it is useless. But

once it is connected to another machine, it becomes useful. So it is with all communication networks, and the TCP/IP Internet is no exception (Sundararajan 2006).

The presence of such contingent motivations for adopting TCP/IP (rather than just for its technological superiority) has led some historians and economists to speculate if TCP/IP won out because of some sort of lock-in effect. Lock-in in this context refers to a situation where a dominant technology becomes so simply because of its earlier entry into a market, and it can possibly “prevent higher quality products from entering the market” (Maathius 2003, 162). Given the implementation issues with OSI (a lack of running, easily compatible standards for years) and the capacity issues with TCP/IP (limits in addressing and routing), it’s clear neither technology was perfect at the time of competition (168-169). This is again a gray area, where the differing values of network builders could influence whether or not they believe decentralized TCP/IP routing is superior.

But lock-in isn’t the only way to examine socially contingent technology choices. Because of the cumbersome bureaucratic processes of ISO, protocols couldn’t be developed any faster... and this leads into the second, more socially based explanation for one of the reasons behind OSI’s failure: Internet builders were afraid of losing their power. In the early 1990s, the IAB began to think about replacing a piece of the Internet Protocol Suite with an OSI protocol, the CLNP (ConnectionLess Network Protocol). As revealed in RFC 1347, the technical reasoning was to replace TCP and UDP (another transport layer protocol) with a protocol with a capability for bigger addresses (Callon 1992). Also, it would somewhat satisfy demands for standardization from the ISO.

But before the proposal went before the IETF community, the IETF participants unceremoniously found out from a press leak and the backlash ensued. It seems there was already some resentment of the top-down approach and formalization embodied in an organization such as ISO, but now IETF contributors felt as if the IAB was disregarding the principles the Internet had been founded on—rough consensus and running code. (Russell 2006, 54-56) In fact, this is the context where this very popular phrase arose. David Clark gave a speech to the hundreds of irate


IETF participants at their 24th meeting in 1994, entitled “A Cloudy Crystal Ball”, that roused and assured the Internet community that the proposal would be withdrawn and the grassroots nature of Internet construction would continue (Denardis 2009, 43-46).

There are other arguments for the victory of TCP/IP, including one for the selfishness of computer manufacturers. Early in the development of OSI, Vint Cerf had submitted TCP as the protocol for the transport layer, but it was rejected because the standard setters were afraid it would give American manufacturers an edge since they had been working with TCP/IP for years (Abbate 2000, 174). Manufacturers obviously went through similar lines of thought. Derek Barber, a researcher at NPL, put it like this: “The upper layers [of the OSI model] were to do with the way computers communicate, and the big manufacturers of computers had their own way of doing it. They didn’t want success and that’s why we never got agreement on the upper four layers” (Gillies 2000, 89). From this perspective, as Denardis put it, TCP/IP might have helped determine which technology companies succeeded.

In December 1991, the IETF offered up an RFC entitled “Towards the Future Internet Architecture” suggesting possible directions for the network’s evolution, recognizing that current trends (increasing network growth) were putting a strain on the fundamental architecture of the Internet. In this same paper, they formally defined the Internet as anything that runs over IP (D. a. Clark 1991). It is a broad and perhaps undiscerning definition to be sure, often excluding email and disconnected components, but it was a striking verbal pronouncement of the technology’s growing momentum. Perhaps Hughes would see this as the beginning of the end for TCP/IP’s younger, more socially constructed days. This next section details more about that growing momentum, as the Internet crossed into commercialism, exploded in size, and arguments over its inner architecture continued.

TCP/IP DURING THE INTERNET EXPLOSION AND COMMERCIALIZATION (1992 – 2000s)

One of the motives for attempting to bring a piece of OSI into the Internet protocols came about in 1990, when network engineers began to realize that the Internet could literally run out of room (Denardis 2009, 140). As explained in the technical portion on TCP/IP, each IP address (running on the fourth version of IP) must be unique, and early designers allocated 32 bits (units of representation) for these addresses. This means that there are 2^{32} possible addresses, or about 4.3 billion, on the current instantiation of IP. (Most people on the Internet today use IPv4, or IP version 4.)



In the 1970s and 1980s, this no doubt seemed like a vast number, but with the number of devices and humans accessing the network exploding, it became clear that soon 4.3 billion would not be enough (142). The graph above shows the explosion of users in just five years, once the Internet backbone was released by NSF to private corporations, and commercial transactions became allowed online. Within half a decade, the number of Internet users had grown from maybe

50 million to 2 billion. It surely is a testament to the scalability of TCP/IP that it was able to handle such an increase.

RFC 1347 proposed that TCP and UDP could be run over CLNP with bigger addresses in June 1992 to combat looming shortages (Callon 1992). The rejection of CLNP was the setting for the sixth version of IP, a new protocol design that was developed by the IETF in the 1990s. But yet again, as Denardis documents in *Protocol Politics*, “protocol selection became power selection” (25). The beginning of the search for a new protocol with more addressing capability was pointed firmly in the direction of IP, at least in name. After the CLNP was roundly rejected by the IETF for its association with an international top-down organization (the ISO), the alternative protocol, not yet decided upon, was named IPng. Such a name was a predictor of who would maintain control, and what politics this replacement would maintain.

Three proposals for IPng were submitted: SIPP, CATNIP, and TUBA (the original CLNP that caused such a fuss in the OSI vs. TCP showdown). An IPng Directorate was assembled that included members from IBM (the computer manufacturer), CERN (a European nuclear research organization), and Cisco (an American network technology corporation) among others (59). It was clear that IPng would not be a piece of the repealed OSI, and meeting minutes from a 1993 ‘IP decide’ meeting shed some light on why it might have been included in the three options in the first place—participants were worried about the possible legal repercussions if they apparently rejected TUBA on nontechnical grounds (53).

Richard Woundy, a network engineer who implemented GOSIP standards on the NSFnet backbone in the early '90s, recalled something of a similar dynamic. He asserts that the GOSIP implementation, at least from his side, was more to show compliance with a policy than anything else (Woundy 2011). Furthermore, RFC 1454 documents how the search for a new protocol had become something of a process to repair the IETF, rather than a search for the technically best protocol (Denardis 2009, 52). T. Dixon, the RFC author, wrote “Despite the few opportunities for major change of this sort within the Internet, the need for speed of development and low risk have led to the proposals being incremental, rather than radical, changes to well-proven existing

technology” (Dixon 1993). This way of operating hearkens back to the idea of ‘rough consensus and running code’, where quick operability might be prized over long-term design.

Recall that connection to the Internet, as defined in RFC 1287, is connection via IP (D. a. Clark 1991). Denardis hints that the momentum of IP, even just in the power of this definitive name, drove the IPng decision. She documents the asymmetry of the decision process—ISO objectives were often painted as politics, and IETF objectives as simply technical, and in reality both objectives had political and technical elements. The SIPP protocol proposal was presented at the next IETF meeting, and officially named IPv6 (60 - 67).

RFC 2460, published in 1998, describes technical aspects of this new standard. IPv6 uses 128-bit addresses, creating room for 340 undecillion (3.4×10^{38}) unique addresses (Deering 1998). Around the same time that IPng was chosen, another network technology was being compared to TCP/IP and creating factions—but the tension was another form of network architecture and control debates. ATM (Asynchronous Transfer Mode) was developed in the 1970s to carry both data and voice messages. It is also a virtual connection technology, similar in relation to the original ARPANET’s NCP and the X.25 interface—it creates unbroken, predetermined paths for packets to travel (Steinberg 1994, 1 - 4). ATM was chosen by the International Telecommunications Union as the standard to carry data and voice traffic at broadband rates.¹³ Telcos framed ATM as a supporting standard for IP (possibly because of the expected IETF reaction to any suggestion of IP replacement so soon after the OSI showdown). Though as a lower level, universally agreed upon standard, ATM could have been viewed as a usurper of IP lowest common denominator importance (3).

The hip technology magazine *Wired* published an in-depth look at the two camps organized for and against ATM in 1996. Those for the technology were called Bellheads and those against were known as Netheads.¹⁴ Their respective beefs against each other were again about the

¹³ Broadband is a term used to describe an Internet connection that has low latency and high data throughput. It is not a precise technical term, but more of a descriptor. A widely cited 2006 OECD report defined broadband as having “download speeds equal to or faster than 256 kbits/s” (OECD Broadband Statistics 2006).

¹⁴ “In broad strokes, Bellheads are the original telephone people. They are the engineers and managers who grew up under the watchful eye of Ma Bell and who continue to abide by Bell System practices out of respect for Her legacy. They believe in solving problems with dependable hardware techniques and in rigorous quality control - ideals that form the basis of our robust phone system and that are incorporated in the ATM protocol. Opposed to the Bellheads

fundamental architecture of the Internet. “At the heart of the ATM debate is really an older argument,” said computer scientist and software developer Brian Reid at the time. “It is the debate between packet-switching fans and circuit-switching fans; two sides with irreconcilably different points of view” (Steinberg 1994, 4).

The dynamics behind ATM vs. IP (as it was no doubt seen by many) were similar to the OSI vs. TCP/IP that they eventually brought about the same result: IP was triumphant, and ATM faded into the background. There were technical reasons—data traffic has grown at much higher rates than voice, to the point of the Internet subsuming voice networks and sending voice traffic over IP (Steinberg, 1997).

Hughes’ concept of technological moment might shed some light on this point in TCP/IP history. The relative ease with which IP won out over ATM suggests this system was maturing, becoming more self-determining. There were perhaps cultural mindsets that played a part—as social groups also mature in the system, they can also contribute to a system’s momentum. Perhaps many influential networkers, believing in the open packet switching network and carrying a fresh mantra of ‘rough consensus and running code’, added to the inertia of TCP/IP by rejecting other forms of organization around the system’s development.

THE MATURATION OF A SYSTEM (TCP/IP TODAY)

Since the late 1990s into the present, TCP/IP and the Internet as a system have seen the rise of technical and social problems—some new, some along similar veins as older issues. The nature of the network, according to certain critics, is in the midst of change that cannot allow the Internet protocols (and the processes that create them) to remain fundamentally the same. Such analyses again correspond with Hughes’ rough pattern of technological system development—systems are made up of various components, and as the system evolves, some of these components can fall behind and require revamping. These components, known as reverse salients, can

are the Netheads, the young Turks who connected the world's computers to form the Internet. These engineers see the telecom industry as one more relic that will be overturned by the march of digital computing. The Netheads believe in intelligent software rather than brute-force hardware, in flexible and adaptive routing instead of fixed traffic control. It is these ideals, after all, that have allowed the Internet to grow so quickly and that are incorporated into IP - the Internet Protocol” (Steinberg, Netheads vs. Bellheads 1994).

sometimes halt a system entirely, and set the stage for the evolution of a new system. Each of the examples that follow—developmental processes, disregard for technical gentlemen’s agreements, and protocol limitations—could be seen as potential reverse salients, and possibly leading to changes in Internet functions.

The Internet development community has seen a lot of growth in the past decade or so. Focusing just on the IETF, its membership saw a ten-fold increase between 1986 and 2001 (Simcoe 2007, 260). Boston University business school professor Timothy Simcoe researched IETF slowdown in 2007, showing that “technical complexity, committee structure, and distributional conflict are correlated with the duration of the IETF standards setting process” (262). Simcoe details corresponding growth in the organization with the commercialization of the Internet, and a majority of those participants holding .com email addresses (denoting commercial affiliation) by 2010, as opposed to previous majorities of .edu or .org (denoting educational or non-commercial organization affiliation). This reflects the overall commercial background of the participants, as opposed to academic or organizational (268). But that wasn’t the only change. Between 1992 and 2000, the median time to develop a standard and publish it as an RFC increased from 198 to 549 days. Simcoe, through statistical analysis, finds three potential explanations: increased technical complexity, growth of IETF in general, and the rise of intra-community conflicts. Though none of these variables is enough to explain the total slowdown on its own, they each contribute to the idea that the IETF’s bottom-up process of ‘rough consensus and running code’ might not be as efficient in a global-wide participatory network (269 - 290).

Such a suggestion, even as it was implied in the architectural baggage of OSI, might become more palatable in the future as emerging Internet technologies begin to take more and more tolls on the existing network. Of course, the limits of TCP/IP in addressing and routing tables (see Footnote 3) have been pointed out, but new uses of the network as well have begun to crowd the wires, and some claim that network resources are being distributed unfairly. TCP represented most of the traffic on the Internet for years, but UDP (the other transport layer protocol that provides less error checking than TCP), has been gaining popularity for its latency-sensitive applications

such as Voice over IP and streaming video (Ross 2000).¹⁵ The problem with uses such as VoIP and streaming video is the amount of bandwidth¹⁶ they consume, creating large costs on the network because of the most prolific users that impact all other users. Such was the thinking behind scaling back techniques that brought the term ‘network neutrality’ into public consciousness.

One consequence of the Internet’s academic origins is the number of gentleman’s agreements embedded in the actual functions of protocols—and the fact that certain users are taking advantage of them. One of the easiest to understand is the anonymous and thus ‘equal’ nature of users online that allowed the proliferation of spam. Spam (unsolicited, usually commercial electronic messages delivered via email) got its start in 1978. A marketing representative for DEC (Digital Equipment Corporation) sent out a mass message on ARPANET to several hundred recipients about a new DEC system (Kleiner 2008). Reaction was swift and negative—users sent multiple complaints to network administrators. The ARPANET was under the Defense Communications Agency’s purview at this time, and so this marketing representative received some strongly worded reprimands. All of his recipients received the message as well—this network was “FOR OFFICIAL U.S. GOVERNMENT BUSINESS ONLY” and “APPROPRIATE ACTION IS BEING TAKEN TO PRECLUDE ITS OCCURRENCE AGAIN” (Templeton n.d.).

What a change from the current state of spam on the Internet today! By 2009, over 90% of all email had become spam of one form or another. Congress passed the CAN-SPAM measure in 2003, which has barely hindered the stream of unwanted messages that takes a toll on network resources as well as Internet users’ time (McMillan 2009).

Spam is possible because of the trusting nature of Internet protocols. Remember that IP headers contain source addresses that tell which two hosts on the edges of the network are communicating. It turns out that this information is easily forged, and since routers are stateless (they don’t keep track of where the packets come from, apart from source addresses), the identity

¹⁵ Because UDP doesn’t continually check if packets arrived between hosts, it can push lots of packets in a continuous stream between hosts more quickly. Commonly, if packets are lost for a few milliseconds, two communicating individuals can just ask for a repeat of the last word. Such speed, if not always perfectly reliable, is good enough for real time viewing or communication (Yoo 2011, 13).

¹⁶ “Bandwidth measures the range of frequencies that the transmission system can carry. In the digital domain, a transmission system with infinite bandwidth would be able to transmit an infinite number of bits per second.” A certain number of bits make up each packet (When 2010, 82).

of email spammers is even more elusive—especially when using botnets.¹⁷ Tracing routes from a bird’s eye view of the network would be one way to rectify the problem of spam, but the choice to locate intelligence on the network nodes and trust users not to abuse anonymity has had security repercussions such as spam.

But intentions don’t have to be malicious to create harm—pure economic rationality can play out against the gentlemanly agreements built into the Net. Two recent examples, cited by Christopher Yoo include opening several TCP connections and the scaling back of bit torrent traffic that landed Comcast in hot water with the Federal Communications Commission (Yoo 53).

Both of these recent examples arose because of the ‘bursty’ nature of traffic on the Internet, and the subsequent issues with traffic congestion and management. The term ‘bursty’ refers to data transfer that isn’t steady and continuous, but characterized by short ‘bursts’ from one host to another (Yoo 40). The term is also used to refer to the pattern of overall traffic on networks—generally, weekday nights and weekend days at certain times (immediately after school, after dinner, etc.) are busiest, leading to the slowest broadband access (Collins 2009). Speed is everything on the Internet—the economic relationships between different network providers and consumers are based on the amount of traffic, its direction, and who can provide the shortest (and thus often fastest) pathways (Kearns 2011). Thus, any sort of maneuvering on the network to increase speed or the amount of bandwidth consumed is an economic, rational endeavor—and the gentlemen’s agreements, the polite end-to-end politics of TCP/IP, are disregarded.

When hosts open multiple simultaneous TCP connections to each other, they consume more available bandwidth than other hosts could be using. Just imagine a set number of pipes, open for transfer first-come, first-serve. If one computer co-opts five of those pipes, those are five pipes that no other computers can use, and that’s potentially five times faster the first computer can send and receive its data. The engineering community set a recommended limit of these simultaneous connections to two for equalizing access and traffic congestion purposes (Fielding 1999). Needless

¹⁷ “The term bot is short for robot. Criminals distribute malicious software (also known as malware) that can turn your computer into a bot. When this happens, your computer can perform automated tasks over the Internet, without you knowing it” (Microsoft n.d.).

to say, this recommendation is routinely unheeded. According to Yoo, first Netscape (a popular web browser in the 1990s) began opening eight connections per host, then other web browsers followed suit. The web browser application Mozilla Firefox permits hosts to open 15 connections, as well as other Internet applications (Yoo 2011, 42).

The congestion created by economically rational actors on the network (and the resulting ‘unequal’ distribution of bandwidth) is one of the reasons that centralized control of the network is still appealing to network managers and providers. In 2008, Comcast (a consumer-facing Internet Service Provider) took action against certain applications exchanging files via Comcast’s managed wires. One of the biggest applications was BitTorrent, a P2P (peer-to-peer) file-sharing application that opens several TCP sessions between hosts to swap various files. This type of traffic was growing to an enormous size, and Comcast decided to do something about it. In 2008, the Associated Press reported (incorrectly, according to Comcast) that Comcast was blocking some of the traffic from the P2P applications (Metz 2007). Comcast claimed that it was simply sending reset packets to certain P2P hosts after their requests to establish connection because many of these users were sending exponentially more traffic than most users, taking an unfair toll of network resources. But the end result was often an inability for BitTorrent users to send their traffic (Eckersley 2007).¹⁸

The ensuing controversy became a part of the ongoing ‘network neutrality’ debates, which have already been alluded to earlier in this paper (see footnote 5) and have become a large political debate on and offline. Network neutrality, a politically charged term, refers to the treatment of packets en route—should network managers be able to distinguish between packets, and treat them accordingly? Those in favor of network neutrality say no—this could destroy the very anonymous, equalizing nature of the Internet that gave rise to the legendary garage start-ups such as Facebook, eBay, and amazon.com. They argue that the broad, ‘open’ platform that is

¹⁸ The RST (reset) option is a collection of bits in the TCP header that, if set correctly, terminate a connection between two hosts. TCP “spoof” packets can be created that match the IP address and port of a communicating computer, and with their reset bits turned on, these spoof packets from a third party can interfere in the communication between two hosts (Eckersley 2007).

the Internet could not exist with differential treatment of traffic. But others, including some network managers say such treatment is just impractical and maybe even untenable—not to mention potentially unfair for users whose low levels of traffic are slowed down by file sharing or other such bandwidth-hungry applications. There is also an argument for fairness and freedom, but in terms of the market and neoliberalism.

The net neutrality debate is an example of how technical, historical and moral arguments are all readily employed in a public setting, as opposed to some of the technical-disguising-political reasoning in the IPng debates that were confined to networking intelligentsia. The net neutrality debate (large-scale, infused with ideological rhetoric) has been described as “a game of political football” by some of the press (N.V. 2010). A recurring issue for the FCC nowadays, net neutrality has become a topic of heated public discussion as well as within the engineering community (Federal Communications Commission 2011).

Net neutrality is particularly interesting from a SCOT or Hughesian perspective because debates about it might represent a turning point for the current system. Are concerns for security and control of the network in a larger social setting enough to bring about revolutionary change? TCP/IP, with its arranged open and trusting politics, might no longer suit the purposes of its users. But at the same time, the debate about net neutrality isn't this clear-cut. There are corporate interests and moral arguments, struggles for control and evocations of possible historical precedents. It's a complex issue because TCP/IP is part of a large system of intersecting social, technological, and political parts—and the debate has been amplified by the global presence of the Internet.

But even though the global aspect of the Internet means more people are paying attention to architectural issues and politics, not every Internet issue is as front-page newsworthy, despite its importance. Turning back to the issue of Internet addressing shortages, one can see another current architectural issue with political underpinnings, but not quite as fraught with the same emotion as net neutrality debates. For over a decade, the IETF (along with other organizations, some international) has been pushing for IPv6 integration. But efforts to move the Internet over to this new standard have been slow at best. The last block of IPv4 addresses was assigned out in

February 2011; most of the Internet still uses IPv4, though efforts to update are intensifying (van Beijnum 2011). Some software and computer manufacturing companies, such as Google, Microsoft, and Facebook have been making products compatible with both versions of IP (Flynn). But the switchover has been in progress for over a decade, and has still largely not happened.

One reason for the slow migration is technical difficulty—IPv6 is not backwards compatible with IPv4 (IPv6 can't work with the output of IPv4) and a large-scale transition could create widespread glitches (Flynn). This was the motivation for the World IPv6 Day on June 8, 2011—a number of high-traffic websites agreed to a trial run of IPv6 in order to get an idea of the scope of the problems that could come with a switchover. The event was deemed a successful trial run of the new standard, but it didn't seem to hurry along adoption to any measurable degree (Internet Society 2011).

University of Pennsylvania professor and network engineer Roch Guerin explained that the reason IPv6 hasn't been deployed on a macro level is because “there is no compelling reason” — at least not immediately compelling (Guerin 2011). So far, the address restraints have been alleviated by short-term solutions such as Network Address Translation (NAT), which acts as an agent for several computers on the same network and allowing them all to use the same address (CISCO 2011). Guerin, along with Comcast and NSF, have been running an IPv6 Monitoring project, and recently published a paper on poor performance of IPv6 routing, and how comparable peer routing quality between IPv6 and IPv4 seems one of the best ways to promote IPv6 adoption (Guerin, IPv6 Adoption Monitor n.d.).

But even in the midst of these technical problems, can political motivations have a hand in the process as well? Denardis compares the relative enthusiasm of other polities such as Japan, China, India and the EU in adopting IPv6 to the slow progress in the United States. Her reasons behind mandates for national upgrades included the fact that the US maintained a relatively large portion of IPv4 addresses (so exhaustion of reserves didn't seem so imminent) and these countries wanted to become “more economically competitive in information technology markets relative to the United States” (110). Such talk of market advantages and the need for government mandates

create some interesting parallels with older TCP/IP events, i.e. when ARPA mandated a TCP/IP switchover in 1983 and the refusal of the ISO to accept TCP as a transport protocol.

The centralized nature of Internet governance in the '80s, when it was a military-funded project for scientists, made widespread coordination a possibility, arguably for the betterment of the network. But the growth of the system, the spread of Internet actors including large companies with their own momentums, and the concerted efforts to make the network as unregulated as possible all combine to make a centralized migration to IPv6 impossible. This current state of affairs is easy to juxtapose in opposition to the centralized network of 1982. As recounted at the beginning of this paper, network managers Vint Cerf and Dan Lynch were able to essentially turn off a big NCP switch for days at a time, pushing operators over to TCP/IP in frustration. There is no such switch or management over today's global Internet, and no community with the kind of sanctioned command that would dare to turn off the Internet to force IPv6 adoption. The Internet is now too important, too decentralized, and too embedded in global politics.

CONCLUSION

If asked to define TCP/IP, an engineer would probably describe layers of abstraction, name some architectural principles (e.g. end-to-end), maybe explain some of the protocols on the layers of abstraction, and perhaps give some background on hardware or software implementation. But it should now be obvious to any reader that TCP/IP is more than some technical rules—it is a complex cultural artifact as well, a system within larger social, technological, and political systems. Given a rudimentary understanding of TCP/IP's functions, one can see the arranged politics (a dumb, trusting middle and smart, anonymous hosts on the edges) maintained by 'rough consensus and running code'. Throughout TCP/IP's history, there have been instances where technologies were tied to particular institutions, and the ensuing struggles for power affected technological trajectories. Finally, using STS methodologies to clarify the social-technological interactions, one can also see that TCP/IP is not an entirely new animal, despite its power in bringing about one of the greatest networks in human history. It has been socially constructed, but it also influences those

who use it through its arranged politics.

For the conclusion of this research paper, I'd like to recall an opening quote from Laura Denardis, which now has content in the previous pages to illuminate her insight: “[*Protocols*] control the global flow of information and make decisions that influence access to knowledge, civil liberties online, innovation policy, national economic competitiveness, national security, and which technology companies will succeed” (6).

GLOSSARY

- ATM: Asynchronous Transfer Mode
- BBN: Bolt, Beranek, and Newman
- GOSIP: Government Open Systems Interconnection Profile
- IAB: Internet Advisory/Activities/Architecture Board
- ICCB: Internet Control Configuration Board
- IETF: Internet Engineering Task Force

- IPS: Internet Protocol Suite
- ISO: International Organization for Standardization
- ISOC: Internet Society
- NCP: Network Control Program
- NWG: Network Working Group
- OSI: Open Systems Interconnection
- PDN: Public Data Network
- RFC: Request for Comment
- STS: Science, Technology, and Society
- TCP/IP: Transfer Control Protocol/Internet Protocol
- UDP: Uniform Data Protocol

BIBLIOGRAPHY

Abbate, Janet. *Inventing the Internet*. Cambridge: The MIT Press, 2000.

Anker, Peter. *Virtual Circuit Switching*. 2005. <http://www.telecomabc.com/v/virtual-circuit.html> (accessed May 3, 2011).

Bijker, Wiebe E., Thomas P. Hughes, Trevor J. Pinch. *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology*. Cambridge, MA: The MIT Press, 1987.

Braden, Robert. "RFC 1122: Requirements for Internet Hosts." *Internet Engineering Task Force*. October N/A, 1989. <http://tools.ietf.org/html/rfc1122> (accessed N/A N/A, 2011).

Braden, Robert, Joyce Reynolds, and Steve Crocker. "30 Years of RFCs." *Internet Engineering*

- Task Force*. April 7, 1999. <http://tools.ietf.org/html/rfc2555> (accessed N/A N/A, 2011).
- Callon, Ross. "TCP and UDP with Bigger Addresses (TUBA)." *IETF.org*. June 1992. <http://tools.ietf.org/html/rfc1347> (accessed October 2011).
- Cerf, Vint, interview by Judy O'Neill. *Excerpts from an Oral History* (April 24, 1990).
- Cerf, Vinton. "How the Internet Came To be." In *The Online User's Encyclopedia*, by Bernard Aboba. Boston: Addison-Wesley, 1993.
- CISCO. "How NAT Works." *CISCO*. March 29, 2011. http://www.cisco.com/en/US/tech/tk648/tk361/technologies_tech_note09186a0080094831.shtml (accessed December 2011).
- Clark, David. "The Design Philosophy of the DARPA Internet Protocols." *ACM SIGCOMM Computer Communication Review*, August 1988: 1 - 12.
- Clark, David, and L. Chapin, V. Cerf, R. Braden, and R. Hobbt. "Towards the Future Internet Architecture." *IETF*. December 1991. <http://www.ietf.org/rfc/rfc1287.txt> (accessed November 2011).
- Collins, Barry. "Sunday Evening--the new web rush hour." *PC Pro*. January 8, 2009. <http://www.pcpro.co.uk/blogs/2009/01/08/sunday-evening-the-new-web-rush-hour/> (accessed December 2011).
- Comer, Douglas. *Internetworking with TCP/IP Vol. 1: Principles, Protocols, and Architecture*. Upper Saddle River: Prentice Hall, 1995.
- Crocker, Steve. "How the Internet Got Its Rules." *New York Times*, April 6, 2009: A29.
- Eckersley, Peter, Fred von Lohmann, and Seth Schoen. "Packet Forgery by ISPs: A Report on the Comcast Affair." *Electronic Frontier Foundation*. November 28, 2007 . <https://www.eff.org/wp/packet-forgery-isps-report-comcast-affair> (accessed December 2011).
- Day, J.D, Zimmerman, H. "The OSI Reference Model." *Proceedings of the IEEE* 71, no. 12 (June 2005): 1334-1340.
- Deering, S., Cisco, R. Hinden, and Nokia. "Internet Protocol, Version 6." *IEFT.org*. December 1998. <http://tools.ietf.org/html/rfc2460> (accessed December 2011).
- Denardis, Laura. *Protocol Politics*. Cambridge, MA: The MIT Press, 2009.
- Dixon, T. "Comparison of Protocols for Next Version of IP." *IETF*. May 1993. <http://tools.ietf.org/html/rfc1454> (accessed November 2011).
- Federal Communications Commission. "Search results: Net Neutrality." 2011.
- Fielding, R. et al. "Hypertext Transfer Protocol." *IETF*. June 1999. <http://www.ietf.org/rfc/rfc2616.txt> (accessed November 2011).
- Galloway, A.R. *Protocol*. Cambridge, MA: The MIT Press, 2004.
- Gillies, James and Rober Cailliau. *How the Web Was Born: The Story of the World Wide Web*. Oxford: Oxford University Press, 2000.
- Goldstein, Martin. *How We Know: An Exploration of the Scientific Process*. New York, NY: Penum Publishing, 278.

- Guerin, Roch. *IPv6 Adoption Monitor*. <http://mnlab-ipv6.seas.upenn.edu:8080/monitor/> (accessed September 2011).
- Guerin, Roch, interview by Rebekah Larsen. *Office Hours Interview* (September 19, 2011).
- Internet Society. *World IPv6 Day*. 2011. <http://www.worldipv6day.org/> (accessed November 2011).
- Humphrys, Mark. "Uses of Networks." *DCU Computing*.
<http://www.computing.dcu.ie/~humphrys/Notes/Networks/intro.html> (accessed October 2011).
- Huston, Geoff. "10 Years Later." *Internet Society*. June 2008.
<http://isoc.org/wp/ispcolumn/files/2008/06/10years.pdf> (accessed 2011).
- Hedrick, Charles. *Introduction to the Internet Protocols*. Rutgers University, Computer Science Facilities Group, 1987.
- Hounshell, David. *From the American System to Mass Production*. Baltimore, Maryland: The Johns Hopkins University Press, 1985.
- Hoffman, P. and S. Harris. "RFC 4677: The Tao of IETF." *Internet Engineering Task Force*. September 2006. <http://www.ietf.org/rfc/rfc4677.txt> (accessed March 2011).
- Hoffman, P. "The Tao of IETF: A Novice's Guide to the Internet Engineering Task Force." *Internet Engineering Task Force*. October 15, 2011. <http://www.ietf.org/tao.html> (accessed October 2011).
- Internet Society. "Who We Are." *Internet Society*. 2011. <http://www.isoc.org/isoc/> (accessed April 2011).
- Jaffe, Justin. "Happy Birthday, Dear Internet." *Wired Magazine*, December 31, 2002.
- Kapor, Mitch. "Architecture is Politics (and Politics is Architecture)." *Mitch Kapor's Blog*. April 23, 2006. <http://blog.kapor.com/index9cd7.html?p=29> (accessed March 2011).
- Kearns, Michael. "Internet Economics." *UPenn Course Lecture: MKSE 112*. Philadelphia, PA, November 18, 2011.
- Kessler, Gary. "An Overview of TCP/IP Protocols and the Internet." *An Overview of TCP/IP Protocols and the Internet*. November 9, 2010. <http://www.garykessler.net/library/tcpip.html> (accessed March 16, 2011).
- Kim, J., Watanabe, T. "Standardization of the Early Internet: A Search for Socio-Institutional Factors." *SIIT*. Boulder, CO: IEEE, 2002. 282-290.
- Kleiner, Kurt. "Happy Spamiversary! Spam Reaches 30." *NewScientist*. April 25, 2008.
<http://www.newscientist.com/article/dn13777-happy-spamiversary-spam-reaches-30.html> (accessed November 2011).
- Kozierok, Charles. *The TCP/IP Guide*. September 10, 2005.
http://www.tcpipguide.com/free/t_TCPIPOverviewandHistory.htm (accessed March 20, 2011).
- Maathius, Ivo and Wim A. Smit. "The Battle Between Standards: TCP/IP vs. OSI: Victory Through Path Dependency or By Quality?" *Third IEEE Conference on Standardization*

- and Innovation in Information Technology*. Delft, The Netherlands: IEEE, 2003. 161-176.
- Mathison, Stuart L. "Telenet Inaugurates Service." *ACM SIGCOMM Computer Communication Review*, 1975: 24- 28.
- McKenzie, A. and D. Walden. "The ARPANET, the Defense Data Network, and the Internet." *Encyclopedia of Telecommunications*, 1990: 341 - 376.
- McKenzie, Alexander. "INWG and the Conception of the Internet: An Eyewitness Account." *IEEE Annals of the History of Computing* 33, no. 1 (January 2011): 66 - 71.
- McMillan, Robert. "90 Percent of Email is Spam, Symantec Says." *Network World*. May 26, 2009. <http://www.networkworld.com/news/2009/052609-90-percent-of-e-mail-is.html> (accessed November 2011).
- Merryman, R. "ARPAWOCKY." *Internet Engineering Task Force*. June 22, 1973. <http://tools.ietf.org/html/rfc527> (accessed April 2011).
- Metz, Cade. "Comcast Throttles BitTorrent Users." *The Register*. August 22, 2007. http://www.theregister.co.uk/2007/08/22/comcast_throttles_bittorrent_users/ (accessed December 2011).
- Microsoft. "What is a botnet?" *Microsoft Safety and Security Center*. <http://www.microsoft.com/security/resources/botnet-what-is.aspx> (accessed December 2011).
- Miniwatts Marketing Group. *Internet Users in the World*. March 31, 2011. <http://www.internetworldstats.com/stats.htm> (accessed April 2011).
- Mitchell, Bradley. "What Is A TCP/IP Routing Table?" *About.com*. 2012. http://compnetworking.about.com/od/hardwarenetworkgear/f/routing_table.htm (accessed December 2011).
- New York Times. *Net Neutrality*. December 22, 2010. http://topics.nytimes.com/topics/reference/timestopics/subjects/n/net_neutrality/index.html (accessed March 2011).
- Nolle, Thomas. "Taming the Wild, Wild Web." *LA Times*, July 26, 2001.
- "OECD Broadband Statistics to December 2006." *Organization for Economic Co-operation and Development*. December 2006. http://www.oecd.org/document/7/0,3746,en_2649_34225_38446855_1_1_1_1,00.html (accessed November 2011).
- Odom, Wendell. "Neat Net Stories From Vint Cerf and Dan Lynch at Interop." *Networked World*. May 11, 2011 . <http://www.networkworld.com/community/blog/neat-net-stories-vint-cerf-and-dan-lynch-inte> (accessed October 2011).
- Padlipsky, M.A. "RFC 871: A Perspective on the ARPANET Reference Model." *Internet Engineering Task Force*. September 1982. <http://tools.ietf.org/html/rfc871> (accessed February 2011).
- Perlman, Radio. *Interconnections*. Boston: Addison-Wesley, 2000.
- Sundararajan, Arun. "Network Effects." *Arun Sundararajan: Stern School, NYU*. 2006. <http://oz.stern.nyu.edu/io/network.html> (accessed December 2011).
- Segaller, Steve. *Nerds 2.0.1: A Brief History of the Internet*. TV BBooks, 1999.

- SIGCOMM. "ACM SIGCOMM Tutorial: A Technical History of the Internet." *University of Texas*. August 31, 1999. <http://www.facebook.com/Amazon.comBooks> (accessed September 18, 2011).
- Simcoe, Timothy. "Delay and de jure standardization: exploring the slowdown in Internet standards development." In *Standards and Public Policy*, by Shane M. and Victor Stango Greenstein, 260-295. Cambridge: Cambridge University Press, 2007.
- Shields, William M. "Theory and Practice in the Study of Technological Systems." *Dissertation to Virginia Polytechnic Institute and State University*. October 2, 2007.
- Steinberg, Steve G. "Bellheads Retreat." *Wired*, November 1997.
- . "Netheads vs. Bellheads." *Wired*, October 1994.
- Steiner, Peter. "On the Internet, No One Knows You're A Dog." *The New Yorker*. *The New Yorker*. New York City, 1993.
- Russell, Andrew L. "Rough Consensus and Running Code and the Internet-OSI Standards War." *Annals of the History of Computing*, 2006: 48 - 61.
- Ross, Keith W. and James F. Kurose. "UDP Checksum." *Compter Networking: A Top Down Approach Featuring the Internet*. 2000. <http://210.43.128.116/jsjwl/nrxx.asp?id=48> (accessed December 2011).
- Templeton, Brad. "Reaction to DEC Spam of 1978." *Templetons*. <http://www.templetons.com/brad/spamreact.html> (accessed November 2011).
- The Economist. "The Difference Engine: Politics and the web." *The Economist*, December 24, 2010.
- Time Magazine. "Technology: Sharing the Computer's Time." *Time Magazine*, November 12, 1965.
- Urban, Gregory. "Modern World and Cultural Background." *UPenn Course Lecture: Anthropology 004*. Philadelphia, PA, April 25, 2011.
- van Beijnum, Iljitsch. "River of IPv4 Addresses Officially Runs Dry." *ARS Technica*. February 2011. <http://arstechnica.com/tech-policy/news/2011/02/river-of-ipv4-addresses-officially-runs-dry.ars> (accessed November 2011).
- Vest, Tom. "Book Review: Protocol Politics." *The Internet Protocol Journal (CISCO)* 12, no. 4 (December 2009).
- Wilson, Michael. *Lord Save Me From the So-Called Smart Networks*. 11 2006, June. <http://www.arl.wustl.edu/~reInventTheNet/?p-119> (accessed September 2011).
- Winner, Langdon. *The Whale and the Reactor*. Chicago: University of Chicago Press, 1986.
- When, Andrew. *Dot-Dash to Dot.com: How Modern Telecommunications Evolved from the Telegraph to the Internet*. New York: Springer, 2010.
- Woundy, Richard, interview by Rebekah Larsen. *VP of Platform Engineering at Comcast* Philadelphia, PA, (October 5, 2011).
- Yoo, Christopher. "Layering and Internet Policy: Analysis and Critique." *Workshop and Lecture Series on Technology: Policy, Law, and Economics*. Zurich: Swiss Federal Institute of Technology,

2011.

Zittrain, Jonathan. *The Future of the Internet*. New Haven: Yale University Press, 2008.