Engineering a Principle: 'End-to-End' in the Design of the Internet

Tarleton Gillespie
Department of Communication, Cornell University

forthcoming, *Social Studies of Science*, v36 n3, June 2006

Abstract

The term 'end-to-end' has become a familiar characterization of the architecture of the Internet, not only in engineering discourse, but in contexts as varied as political manifestos, commercial promotions, and legal arguments. Its ubiquity and opacity cloaks the complexity of the technology it describes, and stands in for a richer controversy about the details of network design.

This essay considers the appearance, in the 1970s, of the term 'end-to-end' in computer science discourse, and how the term became a point of contention within disputes about how to build a packet-switched network. I argue that the resolution of some of those disputes depended on the transformation of the term from descriptor to 'principle'. This transformation attempted to close specific design debates, and, in the process, made the term dramatically more useful in those discourses beyond engineering that eventually took a keen interest in the design of digital communication networks.

The term, drawn from common parlance and given not only meaning but conviction, was shaped and polished so as to be mobile. As such, it actively managed and aligned disparate structural agendas, and has had subtle consequences for how the Internet has been understood, sold, legislated, and even re-designed.

Keywords: End-to-end, Internet, architecture, language, discourse, network, engineering, representation, technology, law, copyright

> Reality is a scarce resource and is, therefore, the site of a constant struggle. If the world is constituted in one way in the service of one set of purposes for one group of people, it is thereby preempted. (Carey 1992: 24)

The Definition of the Internet

In August 2004, the chief copy editor of *Wired News* announced that the online publication, prone to capitalizing on the Internet, would no longer be capitalizing the word 'internet'. He reasoned, 'In the case of internet, web and net, a change in our house style was necessary to put into perspective what the internet is: another medium for delivering and receiving information. That it transformed human communication is beyond dispute. But no more so than moveable type did in its day. Or the radio. Or television' (Long, 2004). According to a periodical long fascinated with technological novelty, the birth of the Internet is over; its life as mundane part of our cultural landscape has begun, and we should mark the occasion in language.

This grammatical gesture should certainly not be taken as a definitive end. Despite the way it joins countless other similar gestures to slowly naturalize the Internet[1] as a sociotechnical artifact, the process of collectively defining the Internet is still underway. While its founding designers and early enthusiasts may have moved on from their initial proclamations, the process by which the Internet settles in as a component of modern socio-political activity continues: in courtrooms, on the floor of Congress, in corporate boardrooms, in programmers' cubicles, in college dorm rooms.

This process of negotiation is an element of the 'interpretive flexibility' described by Pinch and Bijker (1987). Every technology is shaped by a process of social definition, in which those invested in it struggle not only to implement the technology, but also to narrate what it is and what it is for. These negotiations occur in a range of contexts, and each is associated with specific kinds of consequences: a court decision about appropriate Internet use can powerfully articulate what the Internet is, but so can the design of a new peer-to-peer application meant to defy that decision. If we hope to understand the Internet as a social as well as a technical artifact -- as an 'Internet culture'[2] developing around and alongside the Internet itself -- we must certainly consider these disputes closely. They offer vital insights into the social construction of a communication technology: one still in its adolescence, despite the grammatical dictates of online news services.

In this essay I would like to address a dimension of these disputes that often goes unnoticed, and has largely been overlooked in the literature on the social construction of the Internet. This dimension is hinted at, but obscured by, the stylistic change proposed by *Wired News*. While these debates continue, something has indeed been stabilized. Even as the cultural, legal, and economic implications of the Internet remain open for discussion, the Internet is increasingly taken to be a particular material thing with an identifiable, persistent, and describable shape. While computer scientists, engineers, and network operators still argue about how this or that element is or should be designed, for the casual user, the Internet simply *is*.

The fact that the technical infrastructure of the Internet seems increasingly stable is particularly important when it comes to debates about its use and consequences that occur outside of engineering circles. Discussions about appropriate legislation for the Internet, or appropriate new services to market with it, or appropriate ways to use it, partly depend on assumptions about how it works. Much of this negotiation over what the Internet is and should be is conducted by non-engineers: lawyers, politicians, users, educators, parents, and advertisers. Often participants in these debates make claims about how the Internet works as proof that it should be used in a particular way. But how does a non-engineer come to understand the Internet's material workings, to the extent necessary to participate in such negotiations? How do they know how the Internet works, how it does what it purports to do, how it is designed, and what the implications of its network architecture are? They are by no means experts on the technical dimensions of computer networks, yet they argue about the impact and use of a massively intricate technological system. As such, their claims about what the Internet is as a social phenomenon must largely take on faith what the Internet is as a material artifact. And, I will argue, many of these assumptions about the technology are embedded in the language that accompanies it, tokens that stand in for an understanding of the technology itself and are increasingly taken for granted as the technology becomes commonplace.

Consider a brief example: In *AT&T v. City of Portland*, the 9th Circuit Court of Appeals deliberated on whether the government authority in charge of allocating cable franchises in Portland, Oregon, had the right to demand that AT&T / TCI open their cable broadband lines to competing Internet service providers (ISPs). The court decided that cable does not fall under the 'common carrier' rules that govern other telecommunication services, such as telephony. The details of the case are insignificant for our purposes. What is important is the moment in which the court compares the telecom arrangement to the Internet:

> The Internet's protocols themselves manifest a related principle called 'end-to-end': control lies at the ends of the network where the users are, leaving a simple network that is neutral with respect to the data it transmits, like any common carrier. On this rule of the Internet, the codes of the legislator and the programmer agree.[3]

It is a small moment in an arguably small decision, but a revealing one. There is a blithe matter-of-factness to the court's shorthand description of the structure of the technology, positing a facile analogy between telephone and Internet network design. In this two sentence description, one of the most powerful courts in the nation says that the Internet has a certain shape, and lends its authority to the belief that this shape means that it (and systems like it) must be regulated in ways suited to that shape.

The notion that the Internet has an 'end-to-end architecture' is an increasingly familiar description of the structure and design of computer networks, one that has been significant in these broader cultural negotiations about what the Internet is and should become. To some extent it is an 'accurate' portrayal of the architecture of this technology, but in other ways it is not: the extent to which it is 'true' is itself disputed in engineering debates about network design that continue to this day. But, the portrayal is far from neutral. 'End-to-end' and similar characterizations are deliberate glosses on the technology, symbolic representations of the very shape of the thing in question.

Where does a term like 'end-to-end' come from? The phrase obviously has a commonplace meaning, or constellation of meanings. Most familiar perhaps is the colloquial way of emphasizing the particularly large quantity of things by suggesting we 'lay them end-to-end' -- conjuring up visions of a line of somethings, linked-sausage-style, making abstract rings around the planet.[4] The term has also enjoyed more specific technical meanings, not in the context of network design but in building construction: the layout of floorboards, perhaps, or the components of a plumbing system. It can also express the distance traveled over a bounded space ('we walked that trail from end-to-end'), a slightly different connotation that implies not interconnected components, but the complete path they represent.[5] But how did this term come to stand in as an icon for the shape of the Internet, inside of engineering debates and out, such that authoritative claims about the technology's value, purpose, and impact could be built upon it?

In order to reveal the movement and transformation of this term, I collected documents dealing with the architecture of packet-switched networks, technical and otherwise. My interest in this term was first drawn by Lawrence Lessig's use of it in his arguments on the connection between Internet architecture and its subsequent regulation (Lessig, 2001a, 2002). From his work, and from others making similar non-technical arguments about the consequences of the technical architecture, I traced the term back into the engineering literature in which that architecture was being debated. Very often, these non-engineers pointed to a single reference as the origin of the term, or the key statement of its significance: a 1984 paper by Jerome Saltzer, David Reed, and David Clark. Later, I will discuss how and why this article has come to stand as the definitive regarding this term. Inside the engineering debates, I collected over 400 papers from the 1970s, 1980s, and 1990s that used the term 'end-to-end' in reference to network design; I drew these primarily from journals published by the Association for Computing Machinery (ACM) made available in their Digital Library, and the journals and conference proceedings series published by the Institute for Electrical and Electronics Engineers (IEEE) made available in their Xplore database. These collections were chosen not because network design was decided exclusively within these journals, but because they offered a broad enough range of the relevant debates to reveal in a broad sense the ways in which the term was being used. When particular papers in other journals were regularly referenced as important to the network design debate, I added them to the collection. In order to more fully understand the role played by Saltzer, Reed, and Clark's paper, I sought out other papers in which they referenced or defended their original argument, papers in which they were subsequently cited for their contribution, and papers by others that were more skeptical of their argument. Finally, from simple web searches I found examples of the quite common commercial usage of the term; I collected a number of these instances as well, in order to complement the technical, legal, and political uses of the term with examples of its use inside of a promotional discourse.

By considering this diverse, but by no means exhaustive, collection of documents and the way in which this term has been used, defined, and choreographed as part of a technical and political project, I will reveal how 'end-to-end' came to serve as a descriptor of the structure of the Internet. First, I will describe the technical debates about network design in which this term began to appear, characterizing a set of different meanings the term seemed initially to encompass. Then, I will look closely at the Saltzer, Reed, and Clark paper that engaged with these debates and attempted to articulate and champion a design principle that, they argued, could unite a scattered set of design strategies -- and gave that principle the name 'end-to-end'. I will argue that this document reveals a characteristic kind of rhetorical work in engineering, whereby both the materials and the linguistic shape of the artifact are designed together, with the future deployment of both already in mind. Then I will consider a number of contexts in which this term was taken up outside of engineering debates and used to represent the symbolic shape of the Internet, to argue that matching political arrangements should be adopted. Finally, by looking at recent disputes about digital copyright, I will suggest that the conceptual shape of this term and its prevalence as a shorthand for the shape of the Internet has had consequences for the character and outcome of those debates.

At stake, I believe, are the following questions: who designs the 'architecture' of the Internet, in the discursive sense? How does this articulation of the technology happen in and across multiple discourses? How do institutional agendas and tensions shape that process? How does a term like this move and change, and who instigates these maneuvers? What are the consequences of this process for the political, legal, cultural, and technical life of the Internet? And finally, how should we account for the role of language in the process of designing, articulating, and speaking about the structure of a technology?

The Debate over Network Architecture

Designing a complex artifact like a distributed computer network is both a material and a discursive feat; as much as engineers design new material arrangements of circuits and chips and wires, they also design arrangements of words, concepts, and meanings (Agre, 1997: 14). The extent to which their proposals are considered, appreciated, and deployed depends both on the way their material designs offer compelling solutions to apparent problems, and the way their discursive designs offer compelling characterizations of what the tool can accomplish, as well as how, and most importantly, why it can do so.

Through the 1970s and into the 1980s, significant research was underway at numerous computer science labs, corporate R&D departments, and independent research communities, all working towards the development of computer networks based on the concept of packet-switching. Packet-switching, the idea that information can be broken up into pieces, flung across a network, and recombined by the receiver, is often taken colloquially as the key innovation that made the Internet possible.[6] But, as much as it was a solution to a series of network design challenges, it also posed a new array of problems; it was the beginning of a design strategy, but by no means the end. There were several, significant problems of implementation still left to work out: deciding how the network should handle delays, congestion, lost packets, duplicate packets, errors, security, etc. and, more generally, what the network should be responsible for ensuring and where the necessary functions should be located.

It was in the context of these debates that the term 'end-to-end' first appeared and took root as a technical term. In these early days of inter-network design, disputes over these problems began to coalesce around two models; as these models began to take shape, they each offered an increasingly coherent vision for how to solve many or even all of these technical challenges. For instance, any network system must be able to sufficiently ensure the reliability of information transmission. There will inevitably be failures in the system; consequently, mechanisms must be in place that will recognize such failures and be able to recreate the lost data transfer operations, such that the information gets through regardless of the interruption. Previous data transmission networks (including, until quite recently, the U.S. telephone system) solved this problem by building in a number of intermediary points between sender and receiver; each of these points is responsible for the information that reaches it, and temporarily preserves a copy of all data that pass through until their arrival at the next point is assured. This model, in which intermediate points 'maintain state', resolves interruptions and breakdowns by returning to the last point the data had reached and re-sending them from that point forward. In simpler terms, data moving from point A to point Z must pass through each intermediate point (J, K, L, etc.). If there is a loss of data at point N, the system returns to point M and restarts the data transfer from there.

Some believed that a similar model would be appropriate for packet-switching networks; the hardware at each intermediate node in the network could be given the necessary capacity (data storage, applications that can react to interruptions and recreate messages, etc) to perform this function. This requires an 'intelligent' network (Isenberg 1997), in which each of the interior passage points must be designed to have the necessary resources to manage the flow of data. This model generally came to be known as the 'virtual circuits' model, a nod to the circuits used in analog telephony, and was given official life in the x.25 and x.75 proposals made by the Consultative Committee on International Telegraphy and Telephony [CCITT] in 1976 (Abbate, 1999: 148-156).[7]

Others, however, argued that this strategy would not best take advantage of a packet-switching network; as an alternative, they argued for a 'datagram' model of network design. Rather than making the interior passage points responsible for maintaining state and overseeing data transfer, they argued that only the end-points should manage data; the interior points should be 'dumb' terminals that merely pass along any information that happens to arrive at their doorsteps. To assure reliability, one end point sends the data; the other end point, upon receiving those data, sends a message confirming receipt. In the event of a system failure, the origin point would not receive confirmation and, after a predetermined amount of time, would send the data again. A and Z oversee the transfer; J, K, and L just do as they're told; if A never hears back from Z, A starts the whole process over. Proponents argued this 'stupid' network would be more efficient, demanding less of its interior passage points, although retransmission would consume more resources, since data must traverse the entire route again.[8] (For a discussion of the relative merits of these two approaches, see for example Gitman [1976], Coviello [1979], Lavia and Rhynas [1979], Postel [1980], Boggs et. al. [1980], Rybczynski [1980].)

Though couched almost exclusively in technical terms, this was an argument not only about the 'best' solution, but a broader disagreement about what these networks would even be for -- in other words, what criteria define what counts as 'best'. If the computer network is primarily for the transmission of voice or for real-time streaming of data, in which the arrival of data packets must be timely and sequential, or for remote sharing of computer resources in which delays would hold up other users and waste expensive computer processing cycles, the reliability assurances and internal mechanisms of the virtual circuits model offered greater promise. If, on the other hand, the computer network is primarily for the exchange of data and asynchronous communication, where packet order is of less concern and errors can be rectified without consequence, the datagram model seems preferable.

It's worth noting that much of this research on datagram approaches came out of an interconnected web of academic computer science programs (MIT's Laboratory for

Computer Science, the Stanford Research Institute, and UCLA, among others), computer manufacturers (particularly Bolt, Baranek, and Newman and the Network Analysis Corporation), and military R+D divisions (DARPA and the Defense Communications Agency). Most of the support for the virtual circuits architecture came, not surprisingly, from telephony research labs (such as those at Bell and Bell Canada).[9] It should come as little surprise that the telephony corporations tended to point to voice and streaming as the logical uses to which a network would be put, while the computer science programs emphasized asynchronous data exchange and remote computer time-sharing, tending to urge each towards the respective network models. And with the presence of telephony and computer corporations conducting and sponsoring research, this debate was already both implicitly political and quietly commercial: Who would run such networks? This debate can be read as a harbinger of the competition between computer and telephony industries over network dominance that followed.

In a broad sense, the datagram model largely triumphed in the design of the Internet: most of the functions expected of the network are handled by the sender and receiver terminals, and the nodes internal to the network have been largely freed of all functions except receiving and passing along data packets. The Internet does not depend on a dedicated connection in the way that the virtual circuit model proposed. This 'connectionless' design strategy is, however, far from total: some routing and congestion control functions are implemented by the routers. In particular, servers (an element of the network that, depending on how you look at it, is either internal to the network or the terminal points, either an 'end' or not) do the most to ensure reliable data transmission, impose security, and ease network congestion.

The debate is also far from over. Discussions continue about the proper design of computer networks, and proposals reminiscent of the virtual circuits model regularly resurface. Especially as the Internet has shifted from a limited-access research tool to a public and commercial network, new interests and priorities have emerged that make the old compromises of the datagram model look more like 'vexing dysfunctionalities' (David, 2001: 23), and designers have suggested numerous alternatives that depend on asking routers to do much more than send and receive.[10]

It was during these early debates that the term 'end-to-end' began to appear, drawn from its more commonplace meaning to refer to the shape of the computer network or the location of network functions, typically those proposed in the datagram model. And like its more commonplace meanings, the use of the term in these engineering discussions was by no means singular, though its connotations were roughly consistent. For the most part, the term was used to describe features of, or problems with, the network: yet even within this descriptive function, there was significant variety.

For instance, 'end-to-end' was sometimes used simply to indicate that two networks were adjacent, and met at a single junction point that was an end-point for each network. This use was most similar to the term's more general meaning, i.e. sausage links. 'This paper describes a new 24-channel cable carrier system which is end-to-end compatible with the Western Electric N3' (Poling, 1967: 642). This usage was more common to discussions of cable networks, and largely disappeared from discussions of computer networks. (For similar uses, see Cristal [1966]; Getgen [1967].)

Much more common was to use 'end-to-end' to describe the particular route a single packet of data would or might take in a single instance, from start to finish. 'An important performance measure of these networks is end-to-end delay, which is the elapsed time from the arrival of a message at its source node to the successful delivery of this message to its destination' (Wong, 1978: 344). In the early discussions of network design, minimizing delay was of prime importance. What came to be a logical metric for such delays was the time taken by a data packet from start to finish; while the point of packet-switching design is that packets take different routes to their destination, aggregate measures of end-to-end travel time could speak to the network's average speed and capacity. (For similar uses, see Coviello [1979]; Girard and Ouimet [1983]; Lavia and Rhynas [1979]; Morling, Cain, Neri, Longhi-Gelati, and Natali [1983]; Papantoni-Kazakos [1979]; Tobagi, Borgonovo, and Fratta [1983]; Wong, Sauve, and Field [1982]; Yuill and Pickholtz [1982].)

A third descriptive use of the term referred to the network's capacity to handle some functions at the end-points in the system. One common example of this was in discussions about mechanisms for acknowledging the receipt of a complete message. Many proposed that the receiver send an acknowledgement (or 'ack') to the sender, to signal a completed transaction. 'The ETE (*end-to-end*) Ack serves as a flow regulator between origination and destination and as an indicator to the origination node (upon receiving an ETE Ack) that the destination node has correctly received the message' (Gitman, 1976: 1258). Though clearly related to the previous use, the emphasis here is on the agency and responsibility of the end-points rather than the character of the route between them. This meaning was widely used in the network design discussions throughout the 1970s. (For similar uses, see Baccelli, Gelenbe, and Plateau [1984]; Belsnes [1977]; Boggs, Shoch, Taft, and Metcalfe [1980]; Campbell [1978]; Cerf and Khan [1976]; Chan and Georganas [1979]; Chatterjee, Georganas, and Verma [1977]; Davies [1972]; Gien and Zimmerman [1979]; Kleinrock and Kermani [1980]; Postel [1980]; Pouzin [1981]; Voydock and Kent [1983].)

Finally, it is worth noting a fourth use of the term, one that did not begin to appear until the 1980s. As the datagram model of network design championed by computer scientists won out over the virtual circuit model, researchers at several of the corporate telecommunications labs, particularly at Bell, shifted their efforts to developing telephone networks that used digital relays rather than analog switches. There was a great deal of discussion of how to get digital all the way to the 'last mile' -- into the home, so that the features offered by digital technology (speed, tailoring services to particular users, interactive functions) would not be hampered by a lingering analog link. In this research, the term 'end-to-end' resurfaced:

> The world is in the midst of revolutionary changes in the ways in which various types of information are generated, and how they are used. To meet fully the needs of this information age will require corresponding advances in the capacity to store, transfer, process and deliver information in a host of different forms and formats. Many of these advances will come within the framework of the Telecommunications Network of the future. This paper describes the Switched Digital Capability (SDC) which is expected to be a major ingredient in the transition from what to date has been largely a voice network, to a 'hybrid' network supportive of a wide range of voice and 'data' services. SDC will bring end-to-end switched digital connectivity as an important early step in the evolution of the 100-year old 'telephone' network to the integrated services digital network (ISDN) of its second century. (Kelcourse and Siegel 1983: 29)

(For similar uses, see Bhusri [1984]; Decina [1982]; Gruber and Le [1983]; and Rybczynski [1980].) It is worth noting two things: first, the term can just as easily be adopted to describe a network designed according to the virtual circuits logic; it is not inherently matched to one particular design. Second, the term fits quite neatly not only an argument for the design of the telecommunication network, but also a series of claims about social change through technology. This will not be the last time that a congruence of technical and social proclamations built on this term will appear.

'With neither explicit recognition nor much conviction...'

In these debates, researchers adopted a number of strategies to make the most persuasive case for this or that argument. Some argued for their position by proposing specific solutions to specific problems, hoping they would prove compelling. Others proposed models for networks, even designed their own, to illuminate the potential of their design strategy. Finally, a few attempted to articulate a design principle, to convince in the abstract rather than on a case-by-case basis. One particular paper that intervened in these design debates adopted this last strategy, and used ‘end-to-end’ to name the proposed design principle. The work undertaken in that paper helped to make this term a powerful and mobile resource within those debates, and gave it the capacity to move beyond engineering disputes about the Internet as a technical phenomenon, entering into political, legal, and economic disputes about the Internet as a social phenomenon.

In April 1981, at the Second International Conference on Distributed Computing Systems, Jerome Saltzer, David Reed, and David Clark presented a paper entitled ‘End-to-End Arguments in System Design’. The paper was later published in 1984, in the ACM journal *Transactions in Computer Systems*. This was a significant contribution to the discussion, in part because of the profile and credibility of its authors. All three were working at MIT on projects funded by the Department of Defense’s Advanced Research Projects Agency (DARPA): Saltzer and Reed were part of MIT’s Laboratory of Computer Science, while Clark was the Chief Protocol Architect of DARPA’s Internet Project.

Saltzer, Reed, and Clark’s paper is of particular (though by no means singular) importance for the design of the Internet infrastructure, because it articulated, clarified, and helped champion the datagram model; it has been influential in the disciplines of engineering and computer science, becoming one of the ‘greatest hits’[11] regularly presented in most introductory Computer Science courses. More importantly, the paper is central to the rhetorical power of the term ‘end-to-end’, as suggested by its regular citation in engineering and non-engineering discussions as the iconic origin of the end-to-end principle, so often taken to be the structure of the Internet. Many who subsequently cite the paper describe it as a landmark contribution to the debates. In 1996, Brian Carpenter authored ‘RFC 1958: Architectural Principles of the Internet’. This ‘request for comments’, intended as part of the written documentation for the Internet, points to Saltzer, Reed, and Clark’s argument as one of the guiding principles for network design, a principle beloved by a community skeptical of principles:

> Many members of the Internet community would argue that there is no architecture, but only a tradition, which was not written down for the first 25 years (or at least not by the IAB). However, in very general terms, the community believes that the goal is connectivity, the tool is the Internet Protocol, and the intelligence is end to end rather than hidden in the network ... The key to global connectivity is the inter-networking layer. The key to exploiting this layer over diverse hardware providing global connectivity is the ‘end to end argument’. (Carpenter, 1996)

The principle is also referenced in RFC 2775 (Carpenter, 2000) in light of the challenge it has faced during the Internet’s growing pains, and again in RFC 3439 (Bush and Meyer, 2002), an update of RFC 1958, which suggests that the end-to-end principle leads directly to the principle of simplicity in network design. In the National Research Council’s 2001 report *The Internet’s Coming of Age* (NRC 2001), the end-to-end argument is cited under the subheading ‘success by design’. Saltzer, Reed, and Clark’s paper is the only citation for the term, alongside a reference to RFC 1958.[12]

Saltzer, Reed, and Clark argue that, in nearly all cases, it is best to put the functions of the network in the ends rather than in the intermediary points.[13] First, this arrangement of mechanisms is typically more efficient. More importantly, measures performed inside the network are often insufficient; if end-to-end measures are therefore necessary anyway, the internal measures become redundant. For instance, one network might ensure the successful delivery of a message by a 'hop-by-hop' acknowledgement system: every time the data made one step in the journey, the two nodes would check that the data was unchanged; if no nodes fail in this responsibility and every individual data packet is assured across every step of the journey, the network may assume that the entire message was successfully delivered. However, the recipient still needs to be sure that the content of the message is identical to the message sent. The internal measures cannot assure that they are identical, since they only encountered individual packets, and might not recognize that data had been changed -- an end-to-end acknowledgement is still necessary. If that is the case, Saltzer, Reed, and Clark (1984: 278-82) argued, then the network designer could do away with all of the internal acknowledgement mechanisms, since the single 'end-to-end' confirmation will assure the reliable delivery of all the data across the entire journey. They go on to demonstrate this same redundancy in questions of security, performance and congestion, suppression of duplicate messages, and ensuring message order.

Saltzer, Reed, and Clark also note that, if these network functions were located at the ends rather than built into the network, the network would be more amenable to innovation, even to uses and applications unanticipated by network designers. Single users could add capacities to their own terminals, and so long as the network's role was still the transmission of data packets, the internal points in the network could accommodate them. If, on the other hand, functions were built into the network hardware that imposed specific requirements on the data, the entire network would have to be upgraded to accommodate every new use -- which would tend to discourage all but the most essential innovations.

What do Saltzer, Reed, and Clark accomplish in their paper? In one sense, they merely rename the datagram model. At one point, they explicitly translate the existing debate into their preferred terms:

> Much of the debate in the network protocol community over datagrams, virtual circuits, and connectionless protocols is a debate about end-to-end arguments. A modularity argument prizes a reliable, FIFO *(first in, first out)* sequenced, duplicate-suppressed stream of data as a system component that is easy to build on, and that argument favors virtual circuits. The end-to-end argument claims that centrally provided versions of each of those functions will be incomplete for some applications, and those applications will find it easier to build their own version of the functions starting with datagrams. (Saltzer, Reed, and Clark 1984: 286)

Moreover, they make clear that the solutions they propose here are not their own: 'The individual examples of end-to-end arguments cited in this paper are not original; they have accumulated over the years' (Saltzer, Reed, and Clark, 1984: 285), Even the term itself is by no means their own; as I have already suggested, both the common meaning and its specific and multivalent descriptive use in network design were already well established. It had even been used to articulate the datagram model as a normative design strategy.

Saltzer, Reed, and Clark's contribution to the debate is best understood in more rhetorical terms. In their paper, they engineer a principle for network design, and draw upon the explicit meanings and the array of available connotations for 'end-to-end' to characterize that principle. By drawing together and articulating existing design strategies as part of a coherent design principle, Saltzer, Reed, and Clark intervene in the ongoing debate on a more abstract level, building on recognizable engineering problems but answering them with a sweeping claim to solve all engineering problems in the same way.

Consider the introductory paragraph, in which the authors stake out the discursive terrain upon which they hope to intervene, and begin to characterize the principle they will offer:

> Choosing the proper boundaries between functions is perhaps the primary activity of

> the computer system designer. Design principles that provide guidance in this choice of function placement are among the most important tools of a system designer. This paper discusses one class of function placement argument that has been used for many years with neither explicit recognition nor much conviction. However, the emergence of the data communication network as a computer system component has sharpened this line of function placement argument by making more apparent the situations in which and the reasons why it applies. This paper articulates the argument explicitly, so as to examine its nature and to see how general it really is. The argument appeals to application requirements and provides a rationale for moving a function upward in a layered system closer to the application that uses the function. (Saltzer, Reed, and Clark 1984: 277)

The aspiration here is clear: Saltzer, Reed, and Clark aim to articulate 'design principles' by gathering under a single rhetorical umbrella a series of engineering arguments 'with neither explicit recognition nor much conviction'. This characterization is a revealing one; it is the 'recognition' and 'conviction' that they are after. The task then, is simple: 'This paper articulates the argument explicitly', and in doing so 'provides a rationale' for current and future designs. 'end-to-end' is transformed from a collection of designs into a claim, an argument, a premise -- from an idea to an ideal. It is something that can, from this point on, stand alone -- and stand *for*. If accepted as a principle, it can be portable: picked up as a convenient shorthand, easily taught, even offered as a solution to new engineering challenges beyond the design of computer networks.[14]

Saltzer, Reed, and Clark are also careful to position their principle as the capstone to a series of debates and to link it to several other well-regarded solutions. As noted above, they make explicit reference to the debate between virtual circuit and datagram approaches to network architecture, reframing the debate so that their principle responds to it. They further certify their principle by aligning it with other well-received design principles, most notably the idea of layers:

> It is fashionable these days to talk about *layered* communication protocols, but without clearly defined criteria for assigning functions to layers. Such layerings are desirable to enhance modularity. End-to-end arguments may be viewed as part of a set of rational principles for organizing such layered systems. (Saltzer, Reed, and Clark, 1984: 287)

In doing so, they work to craft their term, and the design strategy it represents, as the underlying principle that will both champion a model of network architecture, and unite a series of squabbles into a single-purpose school of design.

At this point, this account may seem like a classic story about controversy and closure: an ongoing knowledge dispute in which both the claims and the purposes of those claims are up for debate. Fought by competing interests with divergent agendas, the controversy cannot be resolved by any sort of appeal to some convenient 'truth' about communication practice or the fundamental 'nature' of technological design. Instead, closure of the controversy is brought about by the clever characterization of the problem in new terms, terms that help gather consensus among enough designers that alternatives begin to fall away, and the underlying priorities of those designers are installed as the normative understanding (Collins, 1985; Shapin and Schaffer, 1985; Pinch and Bijker, 1987; Gieryn, 1999). In Latour's terminology, we might see Saltzer, Reed, and Clark as making themselves into 'obligatory passage points' within the field of network design; to the extent that the 'end-to-end' principle they engineered and championed might settle in as accepted practice, they would become the spokespersons for the principle, gaining status and credibility in the field as the 'fathers' of the network architecture the principle was used to produce (Latour, 1987).

To some extent, this is what Saltzer, Reed, and Clark were up to. Many who subsequently reference the 'end-to-end argument', both inside and outside of engineering discourse, point to Saltzer, Reed, and Clark's article as its origin point, despite the fact that their article admittedly cribbed its claim from others. That previous work is regularly elided in subsequent references. For example, in the National Research Council report, when the end-to-end principle is cited, the footnote to the 1984 paper begins 'This was first expressed in ...' (NRC, 2001: 36).

All three of the paper's authors make regular reference to the argument in their own later work, often speaking as keepers of its true meaning and lamenting the extent to which it is violated.[15] In one instance, nearly twenty years after their first paper, they observe that their argument has become 'part of the vocabulary of network protocol and operating system designers ... a class of system design principles that organize and guide the placement of function within a system ... a strong rationale to justify design choices' (Reed, Saltzer, and Clark, 1998).

But while the design principle is regularly cited and is now a mainstay in computer science curricula, it certainly did not bring closure to the debates it responded to, at least not in any simple sense. Debates concerning the virtual circuit and datagram models for arranging network functions continued through the 1980s, sometimes with explicit reference to Saltzer, Reed, and Clark's articulation of the end-to-end principle,[16] but just as often without. As Tanenbaum and Van Renesse (1985: 430-431) put it, 'It is our observation that both arguments are valid, but, depending on whether one is trying to forge a collection of small computers into a virtual uniprocessor or merely access remote data transparently, one or the other will dominate'. These debates erupted again and again in the 1990s around new proposals, like those for active networking and quality of service.[17]

As a principle, however, the term has been deployed in engineering literature with the 'conviction' that Saltzer, Reed, and Clark hoped to give it. In particular contexts it began to stand in as an emblem for a design strategy, and increasingly as the name for the very shape of network architecture. In some ways, 'end-to-end' is the architecture of the Internet in name only, or in only the most abstract sense. But this is important in and of itself, in that the symbolic shape of the Internet, with its specifically crafted discursive character, has since been taken up in discourses well beyond engineering and computer science, drawing on the connotations already given to it there, and further reworking it so as to have life inside of those political, legal, and commercial disputes within which the increasingly public Internet would soon find itself.

Making It Move

It is clear that the term 'end-to-end' has played an important, though by no means simple, role in the engineering debates over approaches to and purposes of network design. At first a primarily descriptive term, it was transformed into a design principle, though this effort was by no means complete or commanding. But if we stop at this point, we have only half the story. For just as engineers design their material artifacts both to respond to their own concerns and to eventually leave their hands and be used by others, so their language is designed not only to persuade each other but also to leave engineering discourse and participate in wider debates. The deliberate depoliticization of engineering discourse, cleansed of its social ramifications as mere problem-solving, only masks the way in which both artifact and discourse are deliberately built both to grapple with political questions, and to move into other arenas where they may contribute to political projects. The Internet itself is one of those contributions; a term like 'end-to-end' is another.

In legal controversies, commercial advertisements, and political screeds, the term 'end-to-end' has found purchase, as a symbolic representation of the Internet and its functioning. But, though it invokes roughly the same semantic point as in the engineering discourse, the resonance of the term in each context is quite different. The adoption of 'end-to-end' inside of non-technical discourses may seem, at first glance, to be the clumsy appropriation of a precise engineering term by people poorly trained in technical principles, transforming it into a looser, more metaphorical icon. But a closer look at its characterization and strategic deployment, both inside of the engineering discourse and in its later manifestations, suggests that something more subtle and significant is going on. It is the term's mobility -- its capacity to work across domains, and to both maintain a certain meaning *and* adapt to the particular connotations of its new discursive context -- that warrants attention. And it speaks to the discursive work involved in the design of technology and the Internet's social construction as a mobile and powerful artifact.

It is this capacity to move across discourses, to deliver value even as it connects with different elements of a new discourse, and to stay intact despite its semantic fluidity, that I want to consider further. As James Bono notes, technoscientific discourse, while it seems to aspire to be autonomous and insulated, must constantly register with other contemporary discourses; 'by fixing meanings in highly specific, local, though still plastic, ways, the diachronic dimensions of scientific discourse come to constitute a synchronically coherent, if now metaphorically reordered and situated, language' (Bono, 1990: 77). Since this means that scientific discourse is inextricably linked to other discourses, these metaphoric associations must tangle and reconcile with broader social and cultural shifts in meaning and values (Bono, 1990: 80).

Political

As the Internet became a cultural phenomenon, self-appointed techno-futurists like John Perry Barlow, Nicholas Negroponte, and Howard Rheingold began to loudly proclaim the potential impact of the new technology in utopian, technomorphic terms. They imagined a dramatic and unavoidable revolution in social and political organization, one that would be induced by the Internet and World Wide Web, and used the technologies as metaphors for the kinds of social arrangements that they believed would sweep away the dead infrastructure of post-industrial modernism. Technology had political implications, according to the *Wired* crowd, and those implications were undeniably liberatory. It is in these stylized political declarations that the term 'end-to-end', dressed in its techno-populist connotations, first found purchase outside of the specific engineering discourse from which it was spawned.

Barlow begins his online manifesto 'Censorship 2000' with the now ubiquitous mantra credited to John Gilmore, 'the Internet treats censorship as though it were a malfunction and routes around it'. He then wonders if the words still apply:

> Of course, they remain true -- to the extent that cyberspace exists in an end-to-end data cloud through which any packet may travel to its address by multiple, unfiltered routes. But increasingly, those routes are being channeled and filtered, while their origins and ends get monitored and placed under legal constraint.
>
> That the Internet's original absence of predetermined information circuits and central switches had political implications was, for some of its many fathers, a feature with intended social consequences. I once asked one of them if he had simply been thinking of designing a system that couldn't be decapitated by nuclear attack. 'I can't speak for the others', he said, 'but I was thinking of a system that didn't have a head'. (Barlow, 2000)

Once willing to insist, with Gilmore, that the technical design of the Internet guaranteed its ability to elude traditional regulation, Barlow laments with some surprise that such control mechanisms are being deployed by states and corporations: 'I would continue to be sanguine about this global outbreak of legislative and regulatory bit-blockage, maintaining my faith in the ability of the Net to route around it -- and, more important, in the persistent cluelessness of the oppressive -- but I find that the forces of control have become more sophisticated' (Barlow, 2000). The fact that he is surprised that restrictive governments like China, Myanmar, and Kazakhstan are regulating access to the Internet and even to computer hardware, and that even Western nations like Germany, Switzerland, and the U.S. are demanding that ISPs regulate access to particularly offensive materials, is a reminder that he presumed this end-to-end network to be a guarantee of end-user freedom -- whereas its designers may have meant only locating control at the ends, a subtle but important difference. While Barlow does not reveal which 'father' he spoke to, it is important to note that his contrast between free network and controlled network is not a neat match to the principle articulated by Saltzer, Reed, and Clark from which he borrows the term.

As Barlow lists the various ways in which governments and corporations are attempting to impose mechanisms of censorship onto the Internet, he premises his argument on the assumption that the Internet as it was designed -- that is, end-to-end -- was a technical guarantee of individual freedom, a bulwark against censorship. The biggest threat to online speech, as Barlow sees it, is not in aspects of the Internet's current design, but in the effort of states and powerful corporations to alter that architecture, the

> many initiatives that would move us away from the end-to-end, packet-switched model -- the model to which John Gilmore referred in his quote at the beginning of this article -- and toward a circuit-switched network rather like the phone system the Internet is replacing ... the eventual effect likely could be the conversion of the data cloud into a complex of predetermined routes that would be easily monitored and censored either by those who operated them or by the governments within whose jurisdictions they operated. (Barlow, 2000)

Again, the distinction between end-to-end and virtual circuit design, once a question of *where* control should be placed, becomes a political argument about *whether* control should be imposed. In Barlow's essay, there is a neat discursive fit between the populist political arrangements he seeks and the technical design of the network that he believes hands users power -- the one taken to be the model for the other -- and only a change in the architecture will threaten to undercut that freedom.

The bigger point to be made here is the way in which Barlow (and he's by no means alone) talks about the Internet *as* end-to-end; the unquestioning use of the term to simply name the material shape of the network both fuels his political argument, and overlooks the extent to which network design is much more complex than the term suggests. What was a design principle, a characterization of a set of similar technical solutions abstracted from their details to make a point, becomes here the *de facto* name for the Internet's structure. The icon comes to stand in for the technology, and as such serves as a political rallying cry. The Internet *is* end-to-end, until someone makes it otherwise.

Not only is the term a powerful symbol, invoking the kind of libertarian, individual freedom Barlow desires, but it obscures the complexity of the network's design -- a complexity that helps explain the complexity of attempts to regulate it. In another article focused more specifically on copyright, Barlow suggests that ISPs like CompuServe should not be held legally responsible for users uploading copyrighted works without permission, 'since CompuServe cannot practically exercise much control over the flood of bits that passes between its subscribers' (Barlow, 1994). This notion of what an ISP can and cannot do is mistaken, in part because of the tendency to think of it as internal to this end-to-end system, and therefore just dumb routers. In fact, ISPs have a great deal of ability to regulate what their subscribers post;

when it comes to the process of packet-switching, CompuServe is in some ways an endpoint in the network too.

Legal

The broadly political claims of Barlow and others like him point to, but only begin to really engage with, the complex political and legal debates about the place of the Internet in society. Of particular concern has been the application of existing laws to online communication, a concern that has already reached the highest courts in the United States.[18] But even in these increasingly specialized debates, the term 'end-to-end' has proven a compelling and persuasive icon of Internet structure.

Lawrence Lessig has been perhaps the most prominent and provocative voice in the debate about digital copyright, and more generally about how law and policy should apply to networked technologies. In an article for the *Financial Times* he makes a similar point to Barlow's, arguing that:

> The internet was born a 'neutral network', but there are pressures that now threaten that neutrality. As network architects have been arguing since the early 1980s, its essential genius was a design that disables central control. 'Intelligence' in this network is all vested at the 'end' or 'edge' of the internet. The protocols that enable the internet itself are as simple as possible; innovation and creativity come from complexity added at the ends. This 'end-to-end' design made possible an extraordinary range of innovation. When Tim Berners-Lee started to sell the idea of a 'World Wide Web', he did not need to seek the approval of network owners to allow the protocols that built the internet to run. (Lessig, 2002)

Like Barlow, Lessig builds his political argument on the underlying assumption that the Internet *is* end-to-end, indeed was born that way, and that this design ensures an egalitarian neutrality that is only threatened by changes to that design. Drawing explicitly on Saltzer, Reed, and Clark's paper (but emphasizing their argument that end-to-end design facilitates innovation, rather than simply being more efficient) he renames the datagram vs. virtual circuits argument as a contrast between a network 'commons' and a network that is 'owned'.

Again, there tends to be a subtle shift from a question of where control is located to whether it can be imposed at all. Sometimes, Lessig (2000) describes these design strategies as 'principles that govern its control', but at other moments, the rhetorical temptation of positing a free vs. controlled network proves too seductive: 'At the core code layer, the network was free. The end-to-end design assured that no network owner could exercise control over the network' (Lessig, 2001b). It is important to remember that, though the proponents of end-to-end design knew they were implicitly arguing against the likes of Bell Telephone, the principle that control should be built into the ends does not in any way exclude the possibility that the entire network itself might become the domain of, or be 'owned' by, a single institution. The principle applied as much to MIT's internal network as it did to the Internet, according to its proponents. But in Lessig's rhetoric, an end-to-end network is a free one.

He furthers this argument by drawing a connection between this technical structure and a legal structure with similar sociopolitical aspirations:

> We can now see how the end-to-end principle renders the Internet an *innovation commons*, where innovators can develop and deploy new applications or content *without the permission of anyone else*. Because of e2e, no one need register an application with 'the Internet' before it will run; no permission to use the bandwidth is required. Instead, e2e means the network is designed to assure that the network cannot decide which innovations will run. The system is built -- constituted -- to remain open to

> whatever innovation comes along. (Lessig, 2001a: 40)

The 'commons', a longstanding principle inside of property law, represents those material or cultural resources that do not belong to a particular owner and are potentially available to all. The term has taken on renewed importance for intellectual property law and the controversies around copyright, referring not only to unowned work, but those works that once belonged to a copyright owner but have fallen into the public domain after the copyright expired. Here Lessig leans on the rhetorical power of that term, itself an engineered icon inside of legal discourse, to suture together a new metaphor: the 'innovation commons'. With this techno-legal-discursive hybrid he simultaneously describes the Internet and makes a claim for how it should be judged in the eyes of the law (as a shared public good, rather than, say, as a linked array of privately owned goods or spaces).[19] The resonance of 'end-to-end' (or its cute abbreviation 'e2e', itself a means of further preparing the term for travel) offers a compelling way to fortify this claim, calling explicitly on its sense of 'neutrality', 'free'dom, and 'open'ness to superimpose the technological arrangement and his favored legal solution.

'End-to-end' proves a powerful symbolic term, sometimes for the shape of the Internet with the most potential for resisting restriction, sometimes for the principle of non-discrimination that Lessig uses it to defend. He even extends 'end-to-end' metaphorically back to the telephone system, to note that:

> when the United States broke up AT&T in 1984, the resulting companies no longer had the freedom to discriminate against other uses of their lines. And when ISPs sought access to the local Bell lines to enable customers to connect to the Internet, the local Bells were required to grant access equally. This enabled a vigorous competition in Internet access, and this competition meant that the network could not behave strategically against this new technology. In effect, through a competitive market, an end-to-end design was created at the physical layer of the telephone network, which meant that an end-to-end design could be layered on top of that. (Lessig, 2001b)

Here the term is far from a technical description: he is referring to a virtual circuits network, albeit a regulated one. 'End-to-end' here means the principle of a network freed of restrictions, one disabled from discriminating against uses or users.

It is important to note that the term 'end-to-end' brings with it a number of connotations, depending on what is asked of it. These are not inherent to the term itself, but are produced, to the extent that the conceptual shape proposed by 'end-to-end' resonates with other spatial and organizational metaphors in our culture. Certainly, it fits neatly with and reinforces other popular metaphors that have been used to describe the Internet: 'cyberspace', 'information superhighway', 'electronic frontier', 'blogosphere'. But more than that, to the extent that our understanding of hierarchical social arrangements is often represented spatially in terms of centers and peripheries, describing a system as 'end-to-end' tends to invoke an aura of populist participation, democratic egalitarianism, openness and neutrality, fairness and inclusiveness. To the extent that systems of power are often characterized in a spatial sense, where those in the center have power and those at the ends do not, 'end-to-end' seems to whisper a call for the reversal of that kind of power, a refusal of hierarchy and regulation, a freedom from domination.[20] The question of the relative power of endpoints and centerpoints maps well onto the slightly broader question of the relative power between end users and central operators involved in a communication network, and onto the much broader question of the relative power between citizens and mediating institutions in social networks.

Commercial

While the strident hype exemplified by Barlow resonated with 60's countercultural rhetoric,[21] it also dovetailed neatly with the massive investment of dollars and dreams by the dotcoms in the late '90s. So the least surprising, though probably the most dramatic, repurposing of 'end-to-end' is its appearance and circulation in promotional discourse. While Barlow invokes 'end-to-end' to remind us of a flattened technical (and perhaps political) hierarchy, and Lessig calls on it to demand a legal regime that holds the Internet to its potential as a public good, corporations eager to sell us Internet and wireless services have borrowed the term to signal a commitment to comprehensive solutions and extensive customer service. Examples are easy to come by:

> AT&T and Time Warner today announced the formation of a significant strategic relationship that will include a joint venture to offer AT&T-branded cable telephony service to residential and small business customers over Time Warner's existing cable television systems in 33 states ... AT&T Chairman and Chief Executive Officer C. Michael Armstrong said, 'Today's announcement with Time Warner will significantly advance AT&T's ability to offer end-to-end 'any distance' communications services to American consumers and businesses'.[22] [Feb 1, 1999]
>
> . . .
>
> Tune in to Sun Microsystems, Inc. this week at Cable 2000 show (Booth #3765). Sun will demonstrate end-to-end interactive television and broadband Internet solutions to enable cable operators and programming creators to develop, provision and deploy new two-way services and content.[23] [May 8, 2000]
>
> . . .
>
> Ericsson is establishing itself as the leading mobile Internet supplier. Unlike any other vendor, Ericsson is able to combine its capabilities in mobile networks, terminals, 3G technologies, multimedia applications, carrier-class IP backbone systems, and professional services to provide a complete, end-to-end solution for any conceivable mobile application.[24] [Dec 1, 2000]
>
> . . .
>
> Nokia (NYSE:NOK) and IBM (NYSE:IBM) today announced that they will together provide end-to-end secured business applications on advanced mobile handsets targeted at the enterprise market.[25] [Dec 17, 2002]

This use of the term also makes claims about the design of the network; in fact, it speaks of a network that was not even in the purview of the engineers that championed the term. They draw together, or perhaps conflate, two sets of connotations: both the sense of locating power at the ends *a la* Saltzer, Reed, and Clark, and thereby in the hands of users *a la* Barlow and Lessig, and the more commonplace meaning of extensiveness, comprehensiveness, total coverage. Doing so, with a term that also signals Internet, new technology, progress, and so forth, helps sell the technologies and applications these vendors offer.

Curiously, the use of the term inside of this promotional discourse elides a distinct difference between the networks being offered and the design principle articulated as 'end-to-end'. When Clark, writing with Marjory Blumenthal, recently revisited the 'end-to-end' argument in light of technological and institutional changes surrounding the wider uses of the Internet, he accused hardware manufacturers and service providers of forcing Internet design away from their beloved principle. ISPs, in an attempt to differentiate their services from their competitors, want to offer an array of user-friendly bells and whistles, which tend to require that the network be 'aware' of particular users and uses; but 'any action that an ISP undertakes to enhance its role beyond basic packet forwarding is not likely to be compatible with end to end thinking, since the ISP does not have control over the end-points' (Blumenthal &

Clark, 2001: 18). The kind of consumer synergy inside these commercial 'end-to-end' promises is, in Clark's view, precisely antithetical to end-to-end network design. Nevertheless, the term seems able to move into the sales pitches of such providers without stumbling on this contradiction.

The Implications of 'End-to-End': the Copyright Debate

'End-to-end' developed its metaphorical resonance inside of engineering discourse. Not only was it deployed as a conceptual benchmark for a philosophy of network design; it also became a commonplace and unproblematic descriptor in more specific engineering discussions. 'End-to-end' has become a rationale with its own justification, a living category increasingly obvious to others. It has proven itself to be a successful design strategy, but perhaps only to the extent that a number of successful design strategies have been convened under its symbolic umbrella. Yet each success lends the term discursive force, each use of the term refines and polishes it as a persuasive concept.

This is not the kind of phenomenon for which consequences are easy to pinpoint. They are bound to be subtle and correlational, embedded in the very language we use. The impact of the language used to describe a technology will always travel alongside the technology itself; it is a particular challenge to separate the consequences of the terminology from the consequences of the artifact it describes. As Bazerman (1999: 159) puts it: 'If Edison's light and power had not become parts of daily material practice, all the technical, legal, financial, industrial, and public relations bubbles ... would have burst, leaving little meaning, value, or presence ... In an important sense, the material bottom line warrants the meaning of the words. But the meaning, location, and form of the material are, in an equally important sense, mediated by the words. Electricity and incandescent light, when they appeared, became material representations of ideas they were realizing'. It may be a futile exercise to even distinguish between the two, since we never experience either in isolation.

However, I want to highlight certain features of the debates inside of one particular legal controversy -- copyright -- features that are somewhat easier to explain if we assume that the meanings and priorities embedded in 'end-to-end' have been taken to heart. Barlow, Lessig, and the marketing divisions of AT&T, Time Warner, Sun, Ericsson, and Nokia found specific value in the term, for describing the Internet in ways amenable to their own agendas. But many others who are similarly responsible for and influential in debates about the purposes and value of the Internet -- copyright lawyers, industry lobbyists, politicians, activists, users -- seem to have embraced, if not the term itself, at least its particular characterization of the shape and character of the Internet. What follows are merely correlations, which admittedly may or may not suggest causation.

In recent debates about copyright law and the distribution of music and other cultural works online, those on both sides have been strikingly obstinate. Rather than looking for assurances that the legal and technological regulation of control would be built to best preserve certain rights and values, users largely refused the idea of control altogether; at the same time, the music and film industries demanded almost total restriction. According to users, information would be free, and there was nothing government or corporations could do about it; according to the industry, information would be free, so drastic measures had to be taken to avoid total collapse. This ideological chasm proved difficult to bridge, even as particular cases arose. When the RIAA sued the makers of the Napster p2p[26] file-trading system, fans were certain that Napster could not be shut down; as it became clear that it could and would be, they boasted that more decentralized applications like Gnutella and Kazaa could not be shut down. Steps taken by the industry have come as a series of surprises to the p2p faithful, all of whom seemed sure their particular Titanic was unsinkable. On both sides of the divide, the decentralized, egalitarian, 'stupid' network simply meant that the power resided at the ends only. In this case, end-to-end-ness, not the actual design of the network but the iconographic principle it represents and the egalitarian connotations it implies, may have suggested that file-trading would be unstoppable.

Even the culture industry, which obviously had a vested interest in imposing systems of regulation for the online transmission of data they considered to be theirs, were surprisingly slow to pursue the bottleneck tactics that they eventually did. In the years before Napster, as the mp3 compression format took root and digital copies of popular music became increasingly available on the web, the RIAA spent an inordinate amount of its time going directly after end-users, trying to make examples of particularly active providers of music. Only when Napster arrived did they petition the courts to hold intermediaries liable. Now, their tactics include going after those intermediaries: pressuring ISPs and universities to regulate use, prosecuting the designers of tools and services. Even those efforts have been meant primarily to force ISPs to reveal the names of particular users, as part of a renewed effort to use high-profile subpoenas to go after end-users.[27] This strategy came relatively late in the game -- perhaps too late.

The 1998 Digital Millennium Copyright Act (DMCA) gave ISPs, routers, and search engines a clear protection from liability for copyright violations that might have depended on their material resources. If a commercial service provider or a search engine is alerted to the fact that it is hosting or linking to a website that is in violation of copyright, it need only promptly respond to the 'take-down' request in order to avoid even secondary liability. But consider the language used to indicate who enjoys this benefit:

> A service provider shall not be liable ... for infringement of copyright by reason of the provider's transmitting, routing, or providing connections for, material through a system or network controlled or operated by or for the service provider, or by reason of the intermediate and transient storage of that material in the course of such transmitting, routing, or providing connections, if--
>
> (1) the transmission of the material was initiated by or at the direction of a person other than the service provider;
>
> (2) the transmission, routing, provision of connections, or storage is carried out through an automatic technical process without selection of the material by the service provider;
>
> (3) the service provider does not select the recipients of the material except as an automatic response to the request of another person;
>
> (4) no copy of the material made by the service provider in the course of such intermediate or transient storage is maintained on the system or network in a manner ordinarily accessible to anyone other than anticipated recipients, and no such copy is maintained on the system or network in a manner ordinarily accessible to such anticipated recipients for a longer period than is reasonably necessary for the transmission, routing, or provision of connections; and
>
> (5) the material is transmitted through the system or network without modification of its content. (Section 512a, Digital Millennium Copyright Act)

This characterization of the role of intermediaries is primarily built on a history of regulation of telephone networks, where the courts and regulatory agencies generally adopted a convenient distinction between 'content' and 'conduit'. Like the telephone, mechanisms and institutions merely providing a conduit for information cannot be held responsible for that information. It is this distinction that indemnifies the telephone network if I call someone to plan a crime. But the language of the DMCA's limited liability clause also invokes the principle of 'end-to-end', though not in name: the ISP or search engine can only enjoy this limited liability if it merely provides an automatic process of transmission and routing, does not save copies of the transmitted material, and does not modify the material along the way. This is not the actual end-to-end design of network architecture, but is instead the underlying idea of 'end-to-end' written into the code of law. In fact, it may be written back into the technology itself: in light of this law, network engineers who might prefer adding 'virtual circuit' features like active networking or quality of service mechanisms will enjoy substantially less indemnity from DMCA copyright charges than a network that chooses to remain 'stupid', to honor the end-to-end argument.

Conclusion

The struggle over a technology and its social character is fought, then, on both material and symbolic grounds. Those eager to stake a claim for a technology like the Internet aim to wield not only political power over the design process, but also discursive power to circulate compelling metaphors: metaphors that characterize the technology in ways that elevate the uses they prefer and position them best as the keepers of that technology. By highlighting certain features of the technology and obscuring others, they have the power to frame the entire technology in terms of an assumed set of priorities, and to link it not just to a particular agenda, but to quietly align it with other social and political projects.

This points, first, to the need to consider more carefully the relationship between language and technology. First, the dynamics of language in this case are complex. Meaning is both persistent and malleable; interpretation is both polysemic and consequential. We must look at language as it is designed and deployed to intervene in debates about what a technology can and should be. As Nissenbaum (2004: 216) reminds us, 'our concepts are teleological, not only shaping our thoughts and utterances but facilitating, making awkward, or even depriving us of the facility to think and talk about certain things'. Sally Wyatt recognizes the double complexity of metaphoric language applied to technology: 'metaphors may convey something about the future functions and technological configurations of the Internet and they also may reveal the political assumptions and aspirations of those who deploy them' (Wyatt, 2000: 111). 'End-to-end' is emblematic not only of the prevalence of spatial metaphors for the Internet Wyatt is describing, but of language that (over?)emphasizes 'unstructuredness' and a matching freedom of movement, at the very moment that their structural control mechanisms are being installed.

Second, it is crucial that we expand this attention to the agency behind such discourse, to take into consideration the cultural and institutional contexts of its construction and deployment. This agency is too often cloaked by the language itself: 'as metaphors become embedded within discourses and as actors become less reflexive in their use and choice of metaphors, it may appear that the metaphors themselves are becoming active agents carrying with them expectations about the future' (Wyatt, 2000: 111). As Turner insists, discussing the 'frontier' iconography common to Internet discourse:

> ... this rhetoric has emerged less from the mists of literary history than from the deliberate efforts of a particular community of computer manufacturers, software developers, corporate consultants and academics. ... this emerging elite has drawn on the rhetoric of the electronic frontier in order to identify and manage a series of anxieties brought about by broad changes in work practices and personal mobility over the last twenty-five years – changes triggered by and pervasive in the computer and software industries themselves. (Turner, 1999)

By highlighting the players in these developing communities (in the newsgroups of the Well, in the editing rooms of *Wired*, on the board of the Electronic Frontier Foundation) Turner pinpoints the value they have sought in crafting and deploying the 'frontier' metaphor, the work done to manage its impact on engineering, cultural, and regulatory discourse, and the ways in which the metaphor moved from these specific communities into the broader cultural conversation. Similarly, the engineers debating the proper design of computer networks, and the legal and commercial voices discussing its proper implementation, draw upon particular terms and broader discursive frames to move align the technology with their particular purposes. At the same time, we cannot entirely reduce this to the simple, deliberate choices of discrete individuals; as David Edge (1974, 137) expressed it, 'in some important sense, we do not "choose" a successful metaphor -- rather, it chooses us'.

Finally, this requires us to recognize another way in which technical design is a particularly important process and specifically a political one. Not only do designers have a great deal of say about the artifacts they produce, but they also intervene through the language they adopt. The material and symbolic shape of the artifact are designed together, and the political valence of that artifact is already being grappled with, though often cloaked in a discourse that performs its political neutrality. The linguistic barrier between engineering and the broader discourses of politics, law, and commerce is a significant one. Engineering discourse has quite deliberately established a distinction between engineering and all things social and political, as a means to preserve its 'functional autonomy' (Mulkay, 1976) from the messiness of politics and to maintain a convenient distance from the ethical implications of their work. But this does not mean the discursive choices do not also embody political standpoints, even as they obscure them.

If 'end-to-end' was chosen in part because of its egalitarian connotations, then we must recognize engineers as deliberate political actors. Typical of engineering writing, the language in Saltzer, Reed, and Clark's paper is deliberately apolitical, as are the papers they were responding to. As Reed remembers it, he and his co-authors 'attempted to provide the argument in a value neutral way in the paper, so the reader could make his/her own judgment about its relevance to their situation'.[28] However, as most computer scientists will also admit, they were keenly aware of the political ramifications of their argument:

> The functions that are placed according to the end-to-end argument are functions that are inherently cultural or political -- security, reliability, modifiability, congestion control, routing, privacy -- [they] are end-to-end precisely because they involve human needs, human evaluation, and human benefits that must be traded off against each other ... In fact, the idea of pursuing a thing called 'the Internet' (an ur-network-of-networks) was a political choice -- that universal interoperability was achievable and desirable. It's parallel to 'One Europe' or 'World Government', though not the same. The engineers involved were not ignorant of the potential implications at the political level of that choice.[29]

Just as the design was fundamentally though quietly political, so was the symbolic characterization of that design as 'end-to-end'. The term has connotations, and its selection was no accident. The engineers who adopted it chose the term because it was rhetorically persuasive within a technical dispute, even if they were not explicit about why. And to the extent that the term was persuasive inside of the engineering discourse, 'end-to-end' could then resonate in much the same way elsewhere. The effort to make it resonate, or resonate-able, was part of the agenda in the Saltzer, Reed, and Clark paper; they used the term in a way that was already obliquely aware and quietly confronting of its political connotations. Not only what is in a term like 'end-to-end', but also what is not, helps to map the Internet onto a set of political projects that both precede the design of the Internet, draw on it for justification, and carry it forward.

And finally, we should note that the terms chosen by engineers to describe their technologies can drive not only the cultural and discursive life of the technology, but also its subsequent redesign. Discursive tokens like 'end-to-end' are not only adopted as commercial pitches or policy shorthand, but also in the description of the structure of the technology as it is being revised and revamped. Even as we continue to 'design' the Internet as a social, legal, and political phenomenon, we are also redesigning it technologically. With work on a high-speed Internet2 backbone and the recent proliferation of wireless access networks, metaphors for the material shape of the Internet may eventually be *re-materialized* in the Internet itself, porting their connotations into future versions of the technology, and transforming the promise they pursue into a self-fulfilling prophesy.

Tarleton Gillespie is an Assistant Professor in the Department of Communication at Cornell University, and is affiliated with the Department of Science and Technology Studies and the Department of Information Science. His first book, *Technology Rules: Copyright and the Re-alignment of Digital Culture*, will be published by MIT Press in 2006. email: tlg28@cornell.edu

References

Abbate, Janet (1999) *Inventing the Internet* (Cambridge, MA: MIT Press).
Agre, Phil (1997) *Computation and Human Experience* (Cambridge: Cambridge University Press).
Baccelli, Francois, Erol Gelenbe, & Brigitte Plateau (1984) 'An End-to-End Approach to the Resequencing Problem', *Journal for the ACM* 31/3: 474-85.
Barlow, John Perry (2000) 'Censorship 2000', discussion group post, available at http://amsterdam.nettime.org/Lists-Archives/nettime-l-0007/msg00049.html. Last accessed January 13, 2005.
Barlow, John (1994) 'The Economy of Ideas: A Framework for Rethinking Patents and Copyrights in the Digital Age', *Wired*, 2.03: 85-90, 126-9, available at http://www.wired.com/wired/archive/2.03/economy.ideas_pr.html. Last accessed January 13, 2005.
Bazerman, Charles (1999) *The Languages of Edison's Light* (Cambridge, MA: MIT Press).
Belsnes, Dag (1976) 'Single-Message Communication', *IEEE Transactions on Communications* 24/2: 190-4.
Bhusri, Gurcharan (1984) 'Considerations for ISDN Planning and Implementation', *IEEE Communications Magazine* 22/1: 18-32.
Blumenthal, Marjory & David Clark (2001) 'Rethinking the Design of the Internet: The End-to-End Argument vs. The Brave New World', *ACM Transactions on Internet Technology* 1/1: 70-109.
Boggs, David, John Shoch, Edward Taft & Robert Metcalfe (1980) 'Pup: An Internetwork Architecture', *IEEE Transactions on Communications* 28/4: 612-24.
Bono, James (1990) 'Science, Discourse, and Literature: The Role/Rule of Metaphor in Science', in Stuart Peterfreund (ed.), *Literature and Science: Theory and Practice* (Boston: Northeastern University Press): 59-89.
Bush, R., & Meyer, D. (2002) 'RFC 3439: Some Internet Architecture Guidelines and Philosophy', available at http://www.faqs.org/rfcs/rfc3439.html. Last accessed December 5, 2004.
Campbell, Carl (1978) 'Design and Specification of Cryptographic Capabilities', *IEEE Communications Society Magazine*, November: 15-19.
Carey, James W. (1992) *Communication as Culture: Essays on Media and Society* (New York: Routledge).

Carpenter, Brian (1996) 'RFC 1958: Architectural Principles of the Internet', available at http://www.faqs.org/rfcs/rfc1958.html. Last accessed January 13, 2005.
Carpenter, Brian (2000) 'RFC 2775: Internet Transparency', available at http://www.faqs.org/rfcs/rfc2775.html. Last accessed January 13, 2005.
Cerf, Vint & Robert Khan (1974) 'A Protocol for Packet Network Intercommunication', *IEEE Transactions on Communications* 22/5: 637-48.
Chan, Jackson & Nicholas Georganas (1979) 'Dimensions of Message-Switched Computer-Communication Networks with End-to-End Window Flow Control', *Computer Communications (ACM)* 2/4: 161-164.
Chatterjee, A., Nicholas Georganas, & P. K. Verma (1977) 'Analysis of a Packet-Switched Network with End-to-End Congestion Control and Random Routing', *IEEE Transactions on Communications* 25/12: 1485-9.
Collins, Harry (1985) *Changing Order: Replication and Induction in Scientific Practice* (London: Sage Publications).
Coviello, Gino (1979) 'Comparative Discussion of Circuit- vs. Packet-Switched Voice', *IEEE Transactions on Communications* 27/8: 1153-60.
Cristal, Edward (1966) 'Coupled-Transmission-Line Directional Couplers with Coupled Lines of Unequal Characteristic Impedances', *IEEE Transactions on Microwave Theory and Techniques* 14/7: 114-9.
David, Paul (2001) 'The Beginnings and Prospective Ending of "End-to-End": An Evolutionary Perspective on the Internet's Architecture', working paper, available at http://www-econ.stanford.edu/faculty/workp/swp01012.pdf. Last accessed January 13, 2005.
Davies, Donald (1972) 'The Control of Congestion in Packet-Switching Networks', *IEEE Transactions on Communications* 20/3: 546-50.
Decina, Maurizio (1982) 'Managing ISDN through International Standards Activities', *IEEE Communications Magazine* 20/5: 19-25.
Digital Millennium Copyright Act ['DMCA'] (1998); Pub. L. No. 105-304, 112 Stat. 2860 (1998) (codified at 17 U.S.C. 512 and in various sections of Chapter 12 of Title 17 of the U.S.C.).
Edge, David (1974) 'Technological Metaphor and Social Control', *New Literary History* 6/1: 135-47.
Fry, Michael (1992) 'Practical Trade-offs for Open Interconnection', *Proceedings of the 1992 ACM Annual Conference on Communications*: 1-8. Kansas City, MO.
Getgen, Lawrence (1967) 'Design of the 46B Cable Carrier Terminal', *IEEE Transactions on Communication Technology* 15/4: 637-42.
Gien, Michael & Hubert Zimmerman (1979) 'Design Principles for Network Interconnection', in *Proceedings of the Sixth Symposium on Data Communications (ACM)*:109-19. Pacific Grove, CA.
Gieryn, Thomas (1999) *Cultural Boundaries of Science* (Chicago: University of Chicago Press).
Gillespie, Tarleton (2004) 'Copyright and Commerce: The DMCA, Trusted Systems, and the Stabilization of Distribution', *The Information Society* 24/4: 239-54.
Girard, Andre & Ouimet, Yves (1983) 'End-to-End Blocking for Circuit-Switched Networks: Polynomial Algorithms for Some Special Cases', *IEEE Transactions on Communications* 31/12: 1269-73.
Gitman, Israel (1976) 'Comparison of Hop-By-Hop and End-To-End Acknowledgment Schemes in Computer Communication Networks', *IEEE Transactions on Communications* 24/11: 1258-62.
Gruber, John & Nguyen Le (1983) 'Performance Requirements for Integrated Voice/Data Networks', *IEEE Journal on Selected Areas in Communications* 1/6: 981-1005.
Isenberg, David (1997) 'The Rise of the Stupid Network', *Computer Telephony* August: 16-26.
Johns, Adrian (1998) *The Nature of the Book: Print and Knowledge in the Making* (Chicago:

University of Chicago Press).
Kelcourse, Frank and Eugene Siegel (1983) 'Switched Digital Capability: An Overview', *IEEE Communications Magazine*, January: 29-31.
Kleinrock, Leonard & Parviz Kermani (1980) 'Static Flow Control in Store-and-Forward Computer Networks', *IEEE Transactions on Communications* 28/2: 271-9.
Lampson, Butler (1983) 'Hints for Computer System Design', *ACM Operating Systems* 15/5: 33-48.
Lantz, Keith, Judy Edighoffer, & Bruce Hitson (1986) 'Towards a Universal Directory Service', *Operating Systems Review* 20/2: 43-53.
Latour, Bruno (1987) *Science in Action: How to Follow Scientists and Engineers through Society* (Cambridge, MA: Harvard University Press).
Lavia, A. & D. J. Rhynas (1979) 'A Virtual Circuit Based Transaction Service', *ACM Proceeds of the Sixth Symposium on Data Communications*: 69-74.
Lessig, Lawrence (2000) 'It's the Architecture, Mr. Chairman', op-ed version of testimony give 'In the Matter of the AT&T / MediaOne Merger', available online at http://www.lessig.org/content/testimony/cable/Cable.html. Last accessed January 13, 2005.
--- (2001a) The Future of Ideas: The Fate of the Commons in a Connected World (New York: Random House).
--- (2001b) 'The Internet Under Siege', *Foreign Policy*, November 1, available online at http://www.lessig.org/content/columns/foreignpolicy1.pdf. Last accessed January 13, 2005
--- (2002) 'A Threat to Innovation on the Web', *Financial Times* (December 12), available online at http://www.interesting-people.org/archives/interesting-people/200212/msg00053.html. Last accessed January 13, 2005.
Long, Tony (2004) 'It's Just the "internet" Now', *Wired* (August 16), available online at http://www.wired.com/news/culture/0,1284,64596,00.html. Last accessed, January 13, 2005.
Morling, Richard, Gerald Cain, Giovanni Neri, Marco Longhi-Gelati, & Pierantonio Natali (1983) 'Design of a High-speed Word-Switched Transport Station', *IEEE Journal on Selected Areas in Communications* 1/5: 740-50.
Mulkay, Michael (1976) 'Norms and Ideology of Science', *Social Science Information* 15: 637-56.
National Research Council: Committee on the Internet in the Evolving Information Infrastructure, Computer Science and Telecommunications Board (2001) *The Internet's Coming of Age* (Washington D.C.: The National Academies Press).
Nissenbaum, Helen (2004) 'Hackers and the Contested Ontology of Cyberspace', *New Media and Society* 6/2: 195-217.
Papantoni-Kazakos, P. (1979) 'The Potential of End-to-End Observations in Trouble Localization and Quality Control of Network Links', *IEEE Transactions on Communications* 27/1: 16-30.
Pinch, Trevor & Weibe Bijker (1987) 'The Social Construction of Facts and Artifacts: Or, How the Sociology of Science and the Sociology of Technology Might Benefit Each Other', in Weibe Bijker, Thomas P. Hughes & Trevor Pinch (eds), *The Social Construction of Technological Systems* (Cambridge: MIT Press):17-50.
Poling, Philip (1967) 'Engineering and Application Features of the 4161B Cable Carrier System', *IEEE Transactions in Communication Technology* 15/4: 642-7.
Postel, Jon (1980): 'Internetwork Protocol Approaches', *IEEE Transactions in Communications* 28/4: 604-11.
Pouzin, Louis (1981) 'Methods, Tools, and Observations on Flow Control in Packet-Switched Data Networks', *IEEE Transactions on Communications* 29/4: 413-26.
Reed, David (2000) 'The End of the End-to-End Argument', online post, available at http://www.reed.com/dprframeweb/dprframe.asp?section=paper&fn=endofendtoend.html. Last

accessed January 13, 2005.
Reed, David, Jerome Saltzer, & David Clark (1998) 'Active Networking and End-to-End Arguments', online commentary, available at http://web.mit.edu/Saltzer/www/publications/endtoend/ANe2ecomment.html. Last accessed January 13, 2005.
Rybczynski, Antony (1980) 'X.25 Interface and End-to-End Virtual Circuit Service Characteristics', *IEEE Transactions on Communications* 28/4: 500-10.
Saltzer, Jerome, David Reed & David Clark (1984) 'End-to-End Arguments in System Design', *ACM Transactions on Computer Systems* 2/4: 277-88. Available online at http://www.reed.com/Papers/EndtoEnd.html. Last accessed January 13, 2005.
Shapin, Steven & Simon Schaffer (1985) *Leviathan and the Air-Pump: Hobbes, Boyle, and the Experimental Life* (Princeton: Princeton University Press).
Svobodova, Liba (1984) 'File Servers for Network-based Distributed Systems', *ACM Computing Surveys* 16/4: 353-98.
Tanenbaum, Andrew & Robbert Van Renesse (1985) 'Distributed Operating Systems', *ACM Computing Surveys* 17/4: 419-70.
Tennenhouse, David, Jonathan Smith, W. David Sincoskie, David Wetherall, & Gary Minden (1997) 'A Survey of Active Network Research', *IEEE Communications Magazine* 35/1: 80-86.
Tennenhouse, David & David Wetherall (1996) 'Toward an Active Network Architecture', *ACM Computer Communication Review* 26/2: 5-18.
Tobagi, Fouad, Flaminio Borgonovo, & Luigi Fratta (1983) 'Expressnet: A High-Performance Integrated-Services Local Area Network', *IEEE Journal on Selected Areas in Communications* 1/5: 898-913.
Turner, Frederick. (1999) 'Cyberspace as the New Frontier? Mapping the Shifting Boundaries of the Network Society', posted to the Red Rock News Eater Service, available at http://groups.yahoo.com/group/rre/message/1099. Last accessed January 13, 2005.
Turner, Frederick (2002) 'From Counterculture to Cyberculture: How Stewart Brand and the *Whole Earth Catalog* Brought Us *Wired* Magazine' Ph.D. dissertation. Univ. of California, San Diego.
Voydock, Victor & Stephen Kent (1983) 'Security Mechanisms in High-Level Network Protocols', *ACM Computing Surveys* 15/2: 135-71.
Wong, J. W. (1978) 'Queueing Network Modeling of Computer Communication Networks', *ACM Computing Surveys* 10/3: 341-51.
Wong, J. W., Jacques Sauve, & James Field (1982) 'A Study of Fairness in Packet-Switching Networks', *IEEE Transactions on Communications* 30/2: 346-53.
Wyatt, Sally (2000) 'Talking About the Future: Metaphors of the Internet', in Nik Brown, Brian Rappert & Andrew Webster (eds), *Contested Futures: A Sociology of Prospective Techno-Science* (Aldershot: Ashgate Press): 109-26.
Yuill, Stuart & Raymond Pickholtz (1981) 'Performance Modeling for Packet Networks with Satellite Overflow Channels', *IEEE Transactions on Communications* 29/6: 808-15.

Notes

I want to thank in particular Professors Jerome Saltzer, David Reed, and David Clark for the insights they provided into the intellectual context of their paper. I am also grateful for the helpful comments of Fred Schneider, Chandra Mukerji, Fred Turner, Josh Greenberg, Shay David, and especially the patience and support of Lucy Suchman and Michael Lynch.

[1] For the remainder of this essay, I will continue to capitalize 'Internet', despite the suggestion of the *Wired News* copyeditor. I hope that the argument that follows, like others in this field, suggest that this technology is still far from stabilized as a sociocultural phenomenon.

[2] I use the term 'Internet culture' not mean 'cyberculture', with its implicit contrast to or rejection of realspace activity; instead I use it in the way Adrian Johns (1998) uses 'print culture' -- the technology as well as the institutions, conventions, and rules that develop around it and help make it what it is, and give the technological potential its social particularity.

[3] *AT&T v. City of Portland*, 2000 U.S. App. LEXIS 14383, *23 (9th Cir. 22 June, 2000).

[4] The Oxford English Dictionary contains one reference to the phrase in this specific meaning, a comment from Tyndall in 1860 that '39000 waves of red light placed end to end would make up an inch'.

[5] 'The Air Ground-pipe, laid the whole length of the Green-house ... and reaching from end to end' (Evelyn, 1664: cited in the *Oxford English Dictionary*).

[6] See, for example, Abbate (1999: 7-42).

[7] A particularly useful discussion of these two approaches appears in Postel (1980); in some sense, his discussion is parallel to Saltzer, Reed, and Clark's, without the normative claim that one design model is preferable over the other. See also Abbate (1999: 156-61).

[8] I use transmission reliability merely as one example; while it was perhaps the biggest concern in the debate between the virtual circuit and datagram models, a similar discussion developed in regards to management of network congestion and delay, the process of alerting a receiver that data is on its way, and early considerations of security.

[9] Generally, see Abbate (1999).

[10] See Blumenthal and Clark (2001) for an excellent discussion of this changes and the pressure they exert on traditional network design principles.

[11] Prof. Fred Schneider, informal communication, September 2, 2004.

[12] It is important to note that David Clark was a special advisor to the committee that crafted this research document.

[13] In their paper, they describe this principle as applying to 'systems', which in their discourse include 'networks', but also computer systems, operating systems, and the like. Still, nearly all of their examples are of networks, and they speak directly to the debates about network design; in some sense, the expansion of this point to all 'systems' is a way to give it even more footing in the network argument -- a principle that is undeniable because it is in fact bigger than the question of network design.

[14] For a consideration of their argument, see Tanenbaum and Van Renesse (1985).

[15] See Blumenthal and Clark (2001); Reed (2000); Reed, Saltzer, and Clark (1998).

[16] For example: Lampson (1983), Lantz, Edighoffer, and Hitson (1986), and Svobodova (1984), who are generally in agreement with the end-to-end principle; Fry (1992), and Tanenbaum and van Renesse (1985), who are somewhat more critical.

[17] See, generally, NRC (2001: 107-150). See also Tennenhouse and Wetherall (1996); Tennenhouse, Smith, Sincoskie, Wetherall, and Minden (1997); Reed, Saltzer, and Clark (1998).

[18] For more on Internet copyright controversies, see Gillespie (2003).

[19] And, in case this use of the term seems too far removed from the engineering discourse in which it originated to even make the comparison, rest assured that we have not traveled so far; Lessig, in his expert testimony for the Napster case, quotes the *AT&T v. City of Portland* decision in yet another discussion of 'end-to-end', and cites Saltzer, Reed, and Clark's essay as the origin of the term. See

‘Expert Report of Professor Lawrence Lessig Pursuant to Federal Rule of Civil Procedure 26(a)(2)(B)’, June 2000; *A&M Records, Inc. v. Napster, Inc.*, 239 F. 3d 1004 (9th Cir., 2001).

[20] Ironically, the Saltzer, Reed, and Clark piece uses both an ‘end-to-end’ spatial metaphor as well as an ‘upper-lower’ one -- but, in their model the ends reside at the upper level. To the extent that we typically imagine the control of power at a ‘higher’ point than those subject to that power, their use of a ‘high-low’ metaphor disrupts the mapping of their claim onto traditional understandings of control.

[21] see Turner (2002).

[22] AT&T press release, ‘AT&T and Time Warner Form Strategic Relationship to Offer Cable Telephony’, available at http://articles.findarticles.com/p/articles/mi_m0UNZ/is_1999_Feb_8/ai_53732153 last accessed March 2, 2005.

[23] Sun Microsystems press release, ‘Sun Showcases End-To-End Solutions for Building and Deploying Interactive Digital Television and Broadband Services at NCTA’s Cable 2000’, available at http://www.sun.com/smi/Press/sunflash/2000-05/sunflash.20000508.1.html last accessed March 2, 2005.

[24] Ericsson press release, ‘Ericsson mobile Internet network architecture -- end-to-end technology overview’, available at http://www.ericsson.hu/press/eripress_2000_12_01_230.shtml last accessed February 6, 2003.

[25] IBM and Nokia press release, ‘Nokia and IBM to Provide End-To-End Secured Enterprise Business Applications on Advanced Mobile Handsets’, available at http://www.nokiausa.com/about/newsroom/article/1,1046,986,00.html last accessed March 2, 2005.

[26] ‘p2p’, or ‘peer-to-peer’, is itself another token, and a particularly powerful one, in the copyright debate. It borrows much of its discursive force from ‘end-to-end’, adding a further connotation of egalitarian with the word ‘peer’. I intend to discuss this example fully in a more thorough version of this analysis.

[27] See the memorandum opinion in *RIAA v. Verizon*, available at http://www.eff.org/Cases/RIAA_v_Verizon/20030121-riaa-v-verizon-order.pdf.

[28] Prof. David Reed, personal communication, December 2, 2004.

[29] Ibid.