A PARADIGM SHIFT FOR THE STUPID NETWORK

INTERCONNECTING WITH LEGACY NETWORKS

IN THE INTERNET ERA

By Timothy Denton and François D. Ménard

June 15, 2000

MAIN POINTS

The future of the Internet is the concern of this paper. The relationship of the Internet to carrier networks is its subject. We believe that the benefits to society of the Internet will not be fully achieved unless society in general and policy makers in particular act so as to complete the triumph of the Internet. To do this they must understand how the current legacy networks stand in the way of achieving those benefits. When these facts are more broadly understood, it will be possible for society, and the regulators who serve the interests of the public, to devise policies that accelerate the triumph of the Internet model.

The Internet model is already in the process of transforming what people can expect out of the computer-telecommunications system; and regulators have already allowed it to pose a challenge to the telephone and cable systems. The task now is to complete the triumph of the Internet idea.

The benefits to society of the Internet are derived from its basic functions of connecting computers. The Internet is not a thing or even an arrangement of things, so much as the ability to connect computers of different makes, designs, and operating systems, which ability derives from the instructions to machines that we call protocols. These protocols are and remain the products of the human mind, a fact we emphasize, because if the Internet is to work to its maximum social benefit*, conscious decisions need to be made about it.*

The old task of regulators was to constrain the monopoly powers of owners of transmission facilities: this task persists in the era of the Internet. It persists because the Internet is a way for computers to communicate over transmission systems.

The Internet radically devalues the traditional services offered over transmission systems by stripping out the services they offer; it abolishes the ability of the carrier to constrain the network's abilities, upon which constraints the network's services are defined and its pricing is based.

How the Internet accomplishes this devaluation of the transmission systems over which Internet traffic passes will be the subject of further comment in this paper. For our purpose here, it is sufficient to state that the operators of legacy networks understand this fact about the Internet, they see that the value proposition of the Internet strips away their traditional money-making abilities, and they wish to limit the abilities of the Internet by making it conform to the business models they have inherited, which are highly profitable.

Heretofore, all problems of competition in telecommunications have been about the terms upon which telecommunications carriers would achieve their linkages. These have been complicated enough for regulators to understand. They are now going to be much, much more complicated.

To illustrate this concept, we shall have to refer more extensively to the layered software that comprises the Internet and compare it to the telecommunications system. The telecommunications system interconnects at the *data link* layer, one of the five layers that comprise the Internet and the applications that run on it. All of the problems of interconnection and therefore of competition policy in telecommunications, derive from the connection of the *data link* layer. There are now others for regulators to concern themselves with.

As we shall show in the course of this paper, the layered software model that describes how the Internet works also requires a new understanding of how market power can be exercized. Understanding how the Internet differs from the traditional legacy systems helps us to see the incentives that exist for legacy systems to hold back the Internet for the sake of saving obsolete business models.

The paper concludes with three principles to guide regulatory action in an Internet era. While the basic tasks of the *Telecommunications Act* remain, the principles of constraining market power of incumbent carriers must be applied in new areas, because there are new means of practicing unjust discrimination and undue self-preference.

Those principles are:

- facilities-less competition
- equal access for packets
- minding the protocol stack

The paper explains what is meant by those expressions in the context of interconnecting legacy networks to the Internet.

CHAPTER ONE

THE INTERNET AND LEGACY NETWORKS REPRESENT INCOMPATIBLE IDEAS

The Internet is a complete reformulation of communications. Every concept of the old world of telephony was abandoned by the founders of the Internet. There are no calls, nor are there circuits held open for the call. On the Internet, flows of binary data (or digital circuits, as telephony knows them) are broken up into packets, whose purpose is to limit the amount of information which will be sent at any one time. These packets are then addressed and labeled, using the mechanisms of the Internet Protocol (IP), and they are sent out into the Internet. By performing a process of inquiry of the address encoded in the Internet Protocol destination address contained in the beginning of the packet (the packet header), Internet *routers*, which are the telephone switches of the Internet, then forward the packet to the next router that is more closely connected to the IP address to which the packet is addressed.

However, it is important to note that, until now, Internet routers have only been configured to make their decisions based on the IP address of the packet contained inside the header, and never on the contents of the packet. The Internet does not discriminate among the applications for which it is being used.

Intelligence is placed in the signal itself, unlike the telephone system, where the network is intelligent, and the signal is passively shunted through circuits to its destination.

David Isenberg was the first to call the type of network called forth by intelligent terminals the *stupid network*. The model of services that emerges from the Internet is consistent with how one operates and owns the programs on one's computer. One's relationship to the Internet is the same as one's relationship to one's hard drive. You load applications programs onto it. You offer services with it. The carrier acts as the hard drive, delivering bits where you want them to go, but in no way defining the programs that you can offer or use.

As the Internet is to the highway system, so is telephony to the railway. You drive your car, truck, or vehicle of your choosing over the roadbed subject only to the broadest limitations of weight and speed. The ownership of the car is separate from the ownership of the roadbed. No one specifies when you get on or off the highway. By contrast, the owner of the railbed owns the cars which pass over its rails or it has special interconnection arrangements with other railroads to carry traffic. It is not to be wondered that Canadian telecommunications regulation grew out of railway regulation, since telephony runs on such similar concepts. The contrasts between railways and highways do much to explain both the technical differences between the Internet and the legacy telephone system, and to highlight the shift in power that occurred away from railroads and towards automobile manufacturers. Allowed to triumph, the Internet model foretells a similar transfer of economic power away from carriers.

The relationship of the Internet to the legacy networks on which the packets must still travel is the primary subject of this paper. Because we depend on an existing set of signal distribution media, and because these media incorporate different design philosophies, there is a clash of business interests. On the one hand we have the Internet model: open, equal, end-to-end, peer-to-peer communication. On the other we have the models of telephony and the cable companies, which, despite superficial differences, represent a proprietary, master-slave architecture. The Internet expresses the computer model, and it is from the improvements of design and price-performance in this sector that have derived the social, business and intellectual ferment that characterize the age we live in.

Possibly the most important thing to understand is that the barriers to entry for Internet access providers are minute in comparison to those associated with the telephone network. For the cost of an Internet access feed, a couple of personal computers, a couple of modems and a router, one can become today an Internet Service Provider at a total amortized price of less than $5000 per year. By contrast, the barrier to entry for the telephone network is immense, in the order of 100 to 1000 times more expensive. Once this order of magnitude is understood, it will been seen why any attempt to raise the barrier to entry for Internet Service Providers will immediately be felt by the several thousand entrepreneurs who have chosen the Internet as their model for enabling competition in telecommunications.

The economic boom that has accompanied the rise of the Internet is no accident, but relates to the creation of new forms of wealth and the destruction of old ones. This boom, we believe, derives in a large part from the completely open and unspecified nature of communication through the TCP/IP languages[1]. No one has had to answer to a central authority before launching a product on the Internet. The Internet has no central governing idea of what a product or service is. This central concept of the Internet has allowed a phenomenal unleashing of creative talent, not answerable to the owners of telecommunications systems, whether telephone or cable television. This concept is under attack from the owners of transmission systems, as they seek to delay the end of the their centrally-planned business models[2]. The choke-point is local access, as always. While the area is familiar to regulators, the tools of self-preference are new.

---

[1] The seminal work on the end-to-end architecture of the Net is found in J.H. Salzer, D.P. Reed and D.D. Clark, "End-to-end arguments in system design", *ACM Transactions on Computer Systems*, Nov 1984, p. 277-288.

[2] See Innovation, Regulation and the Internet, Lawrence Lessig, The American Prospect, vol. 11, no. 10, March 27-April 10, 2000, at www.prospect.org/archives/VII-10/lessig-l.html and also Lawrence Lessig's comments to the Ninth International World Wide Web Conference, reported in http://www.wired.com/news/print/0,1294,36283,00.html as "Lessig's Lesson: Beware AT&T"

The tools of self-preference derive from the technologies used to generate and move signals. These technologies are new, and poorly understood by regulators. In a software-driven system, signals have their being and are moved according to the instructions embedded in software - protocols. The telephone industry , and in turn cable television, are being revolutionized by the advent of a system of signal distribution that strips out as much intelligence as possible from the network in favour of machines that lie beyond the edge of the network, in the terminal itself. These terminals are of course computers. Standing between these computers and the Internet are the wires in the ground owned by businesses run on different technical principles. It is no wonder that these incumbents are concerned for their future. Their dominance, in the historical sense, is finished. Their goal is to use the regulatory process, and the lag between what is actual and what regulators and politicians understand the word to be, to delay the advent of business models where they suply the bandwidth, but do not define the services.

If the new model can be delayed by offering “Internet access” which is not actually access to the Internet, if by a few lines of code equal access can be denied, and if the regulator of equal access can be stalled in its tracks for years by technical bafflegab, is this not a small price to pay for continuing the old model? Indeed, is it any price at all?

What role for regulation?

Heretofore, those who owned the wires defined the services. With the advent of the Internet, those who own the physical plant no longer define services. The monopoly of telephone companies and cable companies was not over this or that service; it was the ability to define what services could be extracted from their systems that made them so powerful. The Internet destroys this model.

A conscious public policy for the Internet will seek to transfer power away from those who own wires to those who provide services with Internet applications, because this is where the value is added. Until now they were one and the same people. Since the Internet came into being, networks are reduced to transportation facilities for other people's services and signals. In the telephone era, the last monopoly of the legacy networks was the ability to define services. In the Internet era, their remaining monopoly lies in their ability to prevent certain applications from being used on the Internet, by force or by limiting the speed of the Internet.

We must be very careful using the word "services" in relation to the Internet because the word implies a vertical integration of transport with applications. By contrast, the Internet dissociates transport from applications. Thus, using the term "Internet services" misleads one into accepting the mental model of the telephone era, where the transport is indissolubly linked with the application.

The call for de-regulation of incumbent *carriers* is misplaced

Unfortunately, this fact is not yet well understood. And because it is not, owners of facilities are busy claiming that what they need is a free hand from regulation. It is not from them that progress can derive, because their business model consists of defining what services are or can be. For legacy carriers, it is impossible to conceive that transport can ever be dissociated from applications. Putting it another way, technological innovation in industries does not come from those who ship products, but those who invent them, make them and sell them. No policy for the development of electronic commerce should be based on the views of those who stand to lose the most by it - and these are precisely the legacy carriers. The best means society will ever have to protect itself from a new monopoly in telecommunications is to never impede the creation of innovations on the Internet. By the time one attempts to become a monopoly on a certain application, innovation will have made the monopoly obsolete.

Those who fear that the dead hand of regulation will stifle the Internet fear that, in this rapidly developing field, regulators and policy makers had best stand aside and let the market make all the relevant decisions. This is the new siren song of the entrenched telecommunications giants, whose message of late has been: "let us alone; we will not reach the high-bandwidth future unless you regulators allows us to make the investments, and we will not do so unless you leave us alone to have our way with the Internet." Such an argument is plausible only as long as people fail to distinguish the Internet and the applications it enables, on the one hand, from the physical media over which the Internet works, on the other. Once this distinction is understood, the arguments of the telephone and cable incumbents dissolve.

Economic competition requires rules, boundaries, and referees. Rule-based behaviour is the absolute prerequisite for economic competition. In the particular case of the Internet, the argument for deregulation really amounts to deregulation of those who already own the physical apparatus of wires and cables whereby we reach the Internet. The old monopolies ask regulators and governments to look the other way while they assert their power over other players in the Internet.

The purpose of this essay is to show why such a policy is wrong. Just as the Internet destroyed the world of proprietary - and incommunicative - computers, so the Internet promises the destruction of business model of the carriers. Why this is so is the subject of this essay, and to explain it we must turn our attention to the subject of protocols, the instructions that create and drive the Internet.

How regulatory decisions shaped the Internet

Three major decisions shaped the Internet and assisted Internet penetration in North America:

1. the decision to declare computer services "enhanced" and to exclude them from "telecommunications";
2. the decision to have fixed price local calling areas;
3. the decision to classify Internet service providers as customers of the telephone company.

Enhanced/Basic Services Distinction

In a sequence of rulings in the 1970's and 1980's by the FCC in the United States and in Canada by the CRTC, "enhanced" services were excluded from the definition of "telecommunications" services. The need to do so grew out of the challenges posed by computers. When computation costs were high, computer service bureaux centralized computing in remote locations. As the telephone system became digital - computerized itself - means had to be found to limit telecommunications regulation to companies having "market power". As the computer services industry was vigorously competitive, and showed none of the signs of economies of scale and scope characteristic of telecommunications networks, the enhanced-basic division prevented the spread of price and profit regulation to the entire networked world.

Fixed Price Local Calling Areas

Internet penetration is largely a matter of price. The calling characteristics of Internet use are such that a line has to remain open between the terminal and the Internet service provider. North American telephone policy has always provided for the "free" local calling area. The ability to open a circuit through the legacy system to the Internet, and keep it open for a low monthly price, has been a major factor in stimulating Internet usage in North America. Other countries, such as Germany, are now offering Internet access rates on a similar basis.

ISPs as Customers Rather than Carriers

The decision to treat ISPs as customers of the carriers has had important results, chief of which has been that ISPs continue not to be "carriers", and to have rights of access to the carriers' networks on a non-discriminatory basis. This means that carriers cannot give themselves privileges, such as giving their customers access to the Internet through their facilities on terms better than they allow their ISP customers. Though there are many difficulties to work out, the basic decision of treat ISPs as customers will continue to have positive effects in maintaining the distinction between underlying carriers and the service providers whose signals ride on top of those carriers.

The Internet is not a network

The Internet is not a network; it is a set of protocols for allowing machines to communicate. These protocols are instructions to machines, made by human minds to accomplish particular tasks. Two core protocols that cause computers to link are called Transmission Control Protocol over Internet Protocol, or TCP/IP. Others are involved but the principle remains: the Internet is a set of protocols, a *protocol suite*, which allow otherwise incompatible machines to communicate. It is the Esperanto of machine languages.

The genius of the Internet was to allow communication between computers regardless of the make, operating system or design philosophy of the particular computer. This idea, which is now taken for granted, is fundamentally opposed to the previous world of separate networks, proprietary standards, and control by the owners of these networks of what passes along them. In short, the Internet is the death warrant for all previous business models of communications. This fact is more or less understood by incumbent telephone carriers and cable television network owners, whose revenues have been made in a large part on the restrictions they have been able to impose. Governments have used these restrictions inherent in the old pre-Internet systems for their own purposes, creating subsidies to favoured activities or classes of people. So the Internet undermines these as well.

Dealing With Legacy Networks

By understanding the challenge of the Internet to previous legacy networks, it becomes apparent what tasks lie before regulators of these networks.

We consider that the course of history cannot be reversed, or reversed at a price that the public will tolerate. Accordingly, once regulators and governments launched down the path of allowing the Internet to come into being, they cannot reverse its progress. In the main they actively encourage its growth through appropriate policies. Now the Internet has reached a size and importance where the contrast between what it will allow, and what the old networks allow, is plain to discerning observers.

It is our contention that what now faces governments, and society as a whole, is the role that legacy networks will play in the unfolding drama. Will the constrictions of the legacy networks operate to restrict the enormous capacities of the Internet? At the moment, they do. This may seem strange in the light of the great progress that the Internet has allowed and fostered, but the truth of the matter, as we see it, is that the Internet has not yet reached its full potential because legacy networks, with their obsolete design philosophies, still control access to the Internet, and impose limitations on what could otherwise be achieved if the Internet vision were fully realized.

When we use the term "internet vision" we are relying on the theory advanced by David Isenberg in his concept of the "Stupid Network", which we took up in our paper "Netheads versus Bellheads"[3]. The “Stupid Network” is essentially the reworking of the end-to-end architecture argument advanced by Salzer, Reed and Clark fifteen years earlier. The essence of it can be explained simply. Imagine you had a hard drive in your computer that told you what programs could or could not be loaded onto it. The makers of hard drives would be determining the development of computer applications. Rather than being the passive recipient of programs, hard drives would become a restrictive channel that other makers of applications would have to appease or conform to in order to get onto the personal computer. This describes the role that carriers play, and have always played, in telecommunications. As we stated in "Netheads vs Belheads", the last monopoly is not over this or that service, it is the ability to define what "services" are. As we said earlier in this paper, translated into Internet thinking, this last monopoly is currently being exercised. Certain applications cannot be used on the Internet today: video streaming over the Internet, and running any server applications from home, as demonstrated by recent tariff filings of Canadian cable carriers[4].

In the Stupid Network, the carriers carry the bits. That is all. They provide bandwidth and connectivity, and are paid for it. They do not define which applications can be used, nor necessarily develop the applications themselves. People would have the same relationship to a carrier as they do to the hard drive in their personal computers. They load them up, and when they are full, they buy another. The value in this model is added at the terminal, beyond the limits of the network. Telecommunications companies of the old school have seen this vision and naturally reject it with all their considerable might. Newer players have embraced the Internet vision, but they are largely still laying optical fiber in their backbone networks.

The authors of the end to end philosophy of network design[5] recently had their say:

> This idea of ``stupid operating systems", ``stupid networks", and ``stupid processors" hasn't yet had its full run. The telephone company still seems to think that all users want the illusion of a copper pair from the user's house to some ISP point of presence in another city. In its Internet access approach, the cable company places the real network one step closer, but neither architecture is really prepared for the household with three computers and two network access

[3] Available at http://www.tmdenton.com

[4] In their proposed tariff filings for high speed access to facilities of incumbent carriers, the Canadian cable television operators Rogers, Videotron, Shaw and Cogeco are asking that end-users not be allowed to run any form of Internet server applications from home. This is being conducted under the process deriving from CRTC Decision 99-8.

[5] *Active Networking and End-To-End Arguments,* Comment by David P. Reed, Jerome H. Saltzer, and David D. Clark, at < http://web.mit.edu/Saltzer/www/publications/endtoend/ANe2ecomment.html> Last Modified: 03:17pm PDT, May 15, 1998

providers. Politicians want both the cable company and the ISP to filter packets for things children shouldn't see, and the FBI asks them to make copies of specific data streams to simplify wiretapping (these may be examples of non-cooperating end-points). Political arguments aside, even if one accepts these requirements, the corresponding implementation proposals are sometimes stunning in the way they fail to consider end-to-end arguments.

Protocol interfaces are the new bottlenecks

It is our contention that, for this vision to be realized, there must be instructions created for the machines that carry out these tasks. Such protocols embody the ideas of how the system should work. They are as much the creatures of policy as a regulatory decision is. Interfaces between networks are the places where the machine instructions of one network cooperate or not with another. Hence, when the Internet "interfaces" with a legacy network, the design philosophy of one network may not be able to carry through onto the other. In the case of the Internet, the design philosophy is so different from legacy networks that owners of those networks have every means and incentive to block and frustrate the Internet. Indeed, this is exactly what they are doing. Internet interfaces are thus the new bottleneck facilities to which regulators must direct their attention.

CHAPTER 2

HOW THE INTERNET COMPARES TO TELEPHONY

In order to explain how the Internet compares to telephony, it is useful to compare it to the public switched telephone network, and the technology that underlies it, which today is recognized as circuit switching. The concepts of packets and layers are the foundations of the Internet and completely absent from the telephone network.

In the 1970's, the public switched telephone was transformed into a digital network in order to lower the costs, improve its performance and increase its reliability. This computerization replaced analog equipment by computers without changing any of the fundamental design ideas and assumptions of the circuit switched network[6]. They are:

- the system holds open a circuit from end to end continuously throughout the call;
- the characteristics of the network are built around the usage patterns of humans and their hearing abilities;
- intelligence is scarce and is conserved by placing it inside the system, hence the "intelligent network";
- terminals accordingly are stupid, and are slaved to the nearest central office;
- the purpose of the system is to collect revenue; hence usage is tracked where revenue can be collected.

The point to be made about these design concepts is that no fundamental change accompanied digitization. The move from analog to digital changed the machinery but not the design of the public switched telephone system. Equally important, the assumptions of the telephone network remained firmly embedded in a pre-computer era.

We have described elsewhere[7] how the designers of the Internet started from a completely different set of design assumptions. Intelligence would be plentiful, not scarce, and would be located in the terminal - the computer - as much as possible. Traffic would conform to the characteristics of the machine, rather than the human. So instead of holding open a circuit for what - to a machine - is aeons of time when no transmission would take place, traffic would be broken up into packets. The resources of the network would be usable by other packets and users during that time when some computers were silent. This is the basis of the "always on" feature of the Internet.

Packets are assemblages of code, where the sequence of characters is readable by machines, and by means of which instructions are given about the length, destination, and other important facts about the message. The breaking up of the message into packets means that the message is "switched" millions of times, as each packet is sent out on its

---

[6] They are described in David Isenberg's essay, The Rise of the Stupid Network at http://www.isen.com/stupid.html

[7] See "The Distribution of Signals In Cyberspace", at http://www.tmdenton.com, Chapter Two. (September 1998)

path to find its way to the addressee. Packet-routing technology relies on the speed of computer processing to be feasible. It was designed from the beginning for the characteristics and abilities of computers.

The language of code that eventually prevailed was called Transmission Control Protocol/Internet Protocol, or TCP/IP. Its purpose was to permit communication among machines regardless of the internal operating characteristics. It is therefore a kind of Esperanto for machines, a common language. Thus describing the Internet as a “network of networks” is somewhat misleading; it is rather more a language whereby it is possible to link computers. The network of networks flows out of the common language, and not otherwise.

These machine codes, which are referred to quite properly as protocols, are notionally stacked on top of open another. This is a way of saying that the codes perform certain functions and not others. Changes may be made in a protocol at one layer of the stack without necessarily involving changes in other layers. This is consistent with the idea that the Internet was designed to allow applications to run over different internal systems. The segregation of functions into different layers, each of which consists of a protocol or protocols, is a fundamental design feature of the Internet. Most important then, the Internet segregates applications -the things for which we use and value computers - from the underlying functions of transporting signals. The design philosophy of the Internet is therefore at fundamental variance with telephony, because in telephony the "service" is indissolubly linked to the transport of the signal.

It is for this reason that we must be careful when we use the term "services" in connection with the Internet. The term takes on a fundamentally different meaning when we move from the telephony environment to the Internet. In the Internet environment, the transport is neutral as regard the application that might be transported, such as email or web traffic. The contrary is true with telephony.

The telephone network is comprised of telephone switches linked together via telephone lines. The data, which travels on each of these digital telephone lines, represents the digital encoding of human voice at a rate of 64000 bits per second. At that speed the channel is called a DS0[8]. The function of a telephone switch is to set up and tear down DS0 channels and then to time the duration that the channel was kept open. From this information a phone bill is laboriously computed, which is ultimately passed on to the consumer.

---

8 The reason why 64000 values of either ones or zeros per second are necessary to transmit one second of human voice have their roots in the same science which could be used to explain you how music is stored on a digital Compact Disc. A digitization at a rate of 8000 representations per second is necessary to capture the highest pitch of the voice, which is 4000 Hz. The dynamics of human voice needed to be encoded using 8 bits of information in order to yield the same "toll" quality that people were accustomed to in the non-digital version of the PSTN. By multiplying those 8 bits by the rate of 8000 samples per second, we obtain the value of 64000 bits per second, which is the value of the DS0 unit used as the basis of capacity measurement on the PSTN. A link, which is capable of transferring 24 DS0 in parallel, is called a DS-1. The DS-1 is still used today as the principal unit to measure the interconnection capacity between two telephone switches.

Contrary to the PSTN, the Internet is not a "switched" network, but rather a packet-routed network. The Internet uses routers instead of switches to send packets to a particular destination. The difference between a switch and a router lies in the duration of the switching process. A router ultimately ends up making switching decisions for every packet, whereas a telephone switch makes a switching decision once for every call. The switch makes its switching decision based on an instruction coming from a control network (the Signalling System 7 or SS7 network) and the router makes its decision by looking at the headers inside the packet. A switch has to be warned that it is about to receive data from another switch whereas a router is constantly in a position where it is expecting data from other routers.

Feature	***Circuit-Switched***	***Packet-Routed***
Number of applications possible on the network	Only one: opening and closing the circuit	Potentially unlimited. The value is ultimately unknown unless each application is specified in the headers of the packet, which is never the case.
What can be measured (and thus billed for)	The duration that the circuit is kept open	Every piece of information identified in the header of the packet, the size of every packet and the percentage of link utilization
How it is controlled	By an external control network which instructs switches to set up or tears down the circuits	By link-by-link per packet forwarding control devices, i.e. routers
Knowledge by the system of itself (state)	High degree of self-knowledge	Low degree of self-knowledge

Traditionally, telecommunications was a single application, to apply a computer-era term anachronistically. Telecommunications policy makers are thus experts at regulating *single application networks*. Since the advent of packet-networks, the job of the regulator has become much more complicated as networks are no longer limited to one application.

The Layered Model of the Internet

The purpose of the illustration of the seven-layer OSI model is to explain the relationship between computer applications and the physical substrates upon which they ultimately travel in a packet-routed network environment. The illustration used below shows that the developers of TCP/IP decided not to use the seven-layered OSI model but simplified it by compressing the three top OSI layers into one.

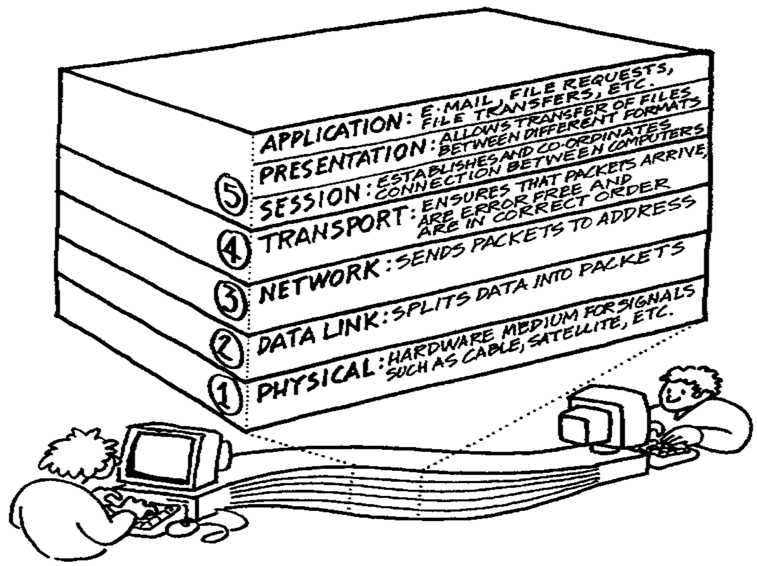

Figure 1. An illustration of the 7-layer OSI model

At the base, telecommunications signals are carried on physical substrates, be they copper, coaxial cable, optical fibre or the air. Telecommunications regulation has always concerned itself with conditions under which customers and competitors would have access to these physical facilities. Generally, competition in telecommunications has been introduced by requiring that incumbent carriers offer their facilities on equal terms to other carriers, and by unbundling elements of the network that were previouslyl tied together. In theory these new entrants could choose the elements of the incumbent's network that were priced right for the competitor, which he would not then have to replace, and to build those elements which the new entrant thought could be more efficiently supplied by his own efforts. This was consistent with the desire to avoid unnecesary duplication of facilities.

Breaking up the elements of a network into several parts and allowing competitors to pick and choose which to rent and which to supply is called *unbundling*.

Unbundling thus happens at Layer 1.

The Relationship of Hosts to Network: Intelligent versus Stupid Networks

The overall architecture of the Internet is best expressed by saying that uses smart hosts (that is, terminals) and a dumb switching network. The idea is that the network should not attempt to impose a restrictive or fixed data flow or connectivity model upon the connected terminals.

The opposite model is that of dumb hosts on a shared smart network: the telephone system. In this case, the hosts require the network to behave predictably under all circumstances, which in turn dictates static partitioning of the network into fixed service elements, where service is either complete or absent. In the case of the PSTN, the service is a 64 Kbps end-to-end circuit established between telephone handsets.

The design philosophy of the Internet is often called an *end-to-end* architecture.[9] Decisions regarding the interaction between the host and the network should, as far as possible, be left to the host (the terminal), rather than the network.

The basic assumption was that the network would undertake best-effort delivery of datagrams (or packets, as they came to be called). Thus any network component may silently discard any packet in response to local congestion. The router is not required to maintain knowledge of the whereabouts of any packet nor maintain a view of the most appropriate path for any packet to follow. The packet itself contains no notion of how it should be delivered, but contains only a global destination identifier, which the local router must translate into a local forwarding decision. To extend a highway metaphor, it is as if an intelligent highway route and direction sign reads the mind of the driver, rather than the driver reading the sign. The car is sent on its way by the sign (the router), rather than by the driver (the packet header) who maintains only the address towards which he is determined to go.

This lack of record keeping within the system indicates that the Internet maintains no "state", to borrow an engineering term. It has very little knowledge of itself. This explains the absence of billing information, which results in the fact that Internet usage is priced on blocks of time connected, rather than on the durations of connections.

Transport Control Protocol (TCP)

[9] See footnote 1 above

TCP is the embodiment of the end-to-end architecture of the Internet. TCP allows data to travel in both directions at once. The packet specifies no predetermined method of data flow control. It allows traffic to build up and when packet loss becomes too great, it backs off until traffic congestion is reduced, and allows traffic to build up again. Hence Internet traffic is adaptive rather than predictive. It sacrifices predictive performance for high utilization and efficiency.

No Central Authority

It can be seen from this description of its basic architecture why the Internet is a self-organizing collection of autonomous networks. No central authority or governing body exists nor can it exist. It is a set of computers that speak to one another in a common protocol. These networks are mostly under private ownership, and the owners have complete authority to shut off communication with other networks as they see fit.

Mathematically its traffic pattern is chaotic, rather than conforming to predictive patterns of the Poisson distribution system that characterizes telephone traffic. This element of chaos is a major asset of the Internet, allowing it to respond to the continuing challenges of scale without the strictures of planning.

Summary of Features

To summarize then, the Internet is characterized by:

1. A layered architecture that separates transport from applications by the TCP/IP layers;
2. An end-to-end architecture that puts the minimum of functions into the network and the maximum of functions into the terminal (computer host) - hence the title " Stupid Network";
3. A chaotic and adaptive routing pattern that sacrifices central control and predictability for maximum network efficiency, which results in ….
4. An absence of central ownership or common planning, as each network is privately owned, and may communicate or pass traffic at the owner's discretion; which results in the fact that…
5. the Internet only knows "autonomous system numbers", that is to say, networks. States and territories have no significance for how it works.

The telephone system is characterized by:

1. A vertical integration of transport functions with service, since the telephony concept predates the possibility of separating "applications" from transport, via computers;
2. A "service" is the product of the intelligence of the network, rather than the user's terminal, hence the "Intelligent Network";
3. A routing system designed around the characteristics of the human voice to produce highly predictable results, based on the calling patterns of humans[10];
4. Central ownership of the network, common planning, and carefully controlled interconnection with other carriers, according to international standards, which leads to…
5. a system based on territories and governments, which work through state-based international agencies, such as the ITU to coordinate standards.

The Internet eliminates vertical feature-integration, or why the telephone system is a *one-trick pony*

We could write a book about these distinctions. For the purposes of this document however, we shall limit ourselves to discussing the first - vertical feature integration of the PSTN, because it is vital to understanding the difference between services in a telephone system, and "services" in an Internet system. While the concept of services is applicable to what we derive from a telephone system, we have to be careful about using the term as if it applied in the same way to the Internet.

Putting it simply, the Internet is a system that supports a multitude of *applications*. These applications may constitute or support *services*. The telephone system supports only one application: telephony. We call it the one-trick pony.

The reason why we refer to the PSTN as a single-application network is that the infrastructure was built for only one purpose: switching 64000 bits per second channels off and on, on demand.

The PSTN is said to use *a vertical integration model* for defining services. Since the process of establishing a call is controlled by the network instead of by the end-user, the network is always aware of the state of the call. This is why we refer to the PSTN as a connection-oriented network. Essentially, the calling model of the PSTN requires that a call be established before applications can reach the end-user.

[10] For example, if everybody simultaneously picked up their telephone, not everybody would get dial-tone.

By contrast the Internet is a connection-less network. At no time is the network aware of the state of applications between two computers, which are interconnected through the Internet. However, for the Internet to be indifferent to the applications that run on it, the architecture of interconnections must be free from self-preference and anti-competitive discrimination. As we have seen throughout the world, owners of legacy systems have the means, motive and opportunity to engage in self-preferential and discriminatory interconnection policy. What is to be done? Chapter Three addresses this issue.

Chapter Three

Basic Ideas for Regulating in an Internet Era

We propose that regulators be guided by three basic ideas in dealing with issues arising in the Internet era. We recognize that in Canada regulators apply the *Telecommunications Act*, and are bound by its provisions regarding unjust or undue self-preference and absence of discrimination, and the limitation of the Act to telecommunications facilities providers.

Remarkably similar concepts underlie the legislation of other countries that regulate telecommunications carriers, since the problem to be addressed derives from the common characteristics of telephone technology.

In essence, the Internet has made the new situation much more complex, as protocols are stacked on top of one another. Opportunities for self-preference by the owner of the physical facilities have been multiplied.

We consider that three principles should guide the regulator in its discharge of its duties to introduce new and relevant forms of competition in previously monopolistic markets. They are:

1. facilities-less competition,
2. equal access for packets, and
3. minding the protocol stack.

The rest of this chapter is devoted to explaining what these principles may mean. It must be emphasized that these are principles to guide action, rather than rules to which everything must be made to conform.

Facilities-less competition

As a general rule, facilities-less competition means that whenever there is a choice to use less or fewer facilities to enable equally effective competition, the bias shall be towards any outcome which results in less or fewer facilities.

Equal Access for Packets

Equal access for packets is a governing engineering principle by which no one is allowed to exercise self-preference by any means. This might include any of the following: to configure, locate, load-balance, render redundant, cache, prioritize or manipulate by any apparatus or method, directly or indirectly, packets originating from a customer and intended for a competitor.

Minding the protocol stack

Minding the protocol stack is a governing regulatory principle that recognizes that opportunities for self-preference and undue or unfair discrimination can arise at higher levels in the protocol stack. Accordingly, regulators must take account of the arrangements of incumbents with their competitors at layers *above that of the physical connection of devices,* and explore the nature of the software by which the machines operate.

Chapter Four explores these ideas in greater depth.

Chapter Four

A Framework for Competition in Wireline High-Speed Internet Access

1. Facilities-less competition

Facilities-*based* competition has been the goal of telecommunications policy in recent years, as regulators have pushed for competition in the telecommunications industry. Telecommunications policy makers came to the conclusion that effective competition could only come from parallel infrastructures. Possibly the best example is the one of the cellular telephony market, where available spectrum was assigned to two or more players. However, it is difficult to conclude that market pricing in a duopoly environment is indeed based on pure survivability of business models in a free market.

Today, the competition for high-speed Internet access to the home is controlled by a wireline duopoly formed by the incumbent telephone carrier and the incumbent cable television carrier. The opportunity to vertically integrate services in an Internet environment can result in far more discrimination than is possible in telephony networks and for this reason, competition based on a duopoly will not be sufficient.

More importantly, if competition is conceived as being based on a duopoly of physical infrastructure providers, then regulators have accepted the obsolete model that vertically integrates services with transport.

Facilities-based competition lies at the very heart of CRTC Telecom Decision 97-8, which introduced local competition in telephony in Canada.

The following quote is from this very decision, at paragraph 73:

> *The Commission is of the view that efficient and effective competition will be best achieved through facilities-based competitive service providers; otherwise, competition will only develop at the retail level, with the ILECs retaining monopoly control of wholesale level distribution.*[11]

Interestingly enough, we find that such a strong statement in favor of facilities-

[11] http://www.crtc.gc.ca/archive/Decisions/1997/DT97-8.htm

based competition is not subject of any elaboration. The intended definition takes form shortly thereafter, when the notions of "essential facility" and mandatory-unbundling are introduced. We argue that the implementation of facilities-based competition through mechanisms of price setting and mandatory-unbundling is tantamount to a concept that we are calling *facilities-less* competition. In short, the CRTC has been pursuing the correct policy.

Our proposed definition of facilities-less competition is the following:

> *Facilities-less* competition means that whenever there is a choice to use less or fewer facilities to enable equally effective competition, the bias shall be towards any outcome which results in less or fewer facilities.

Given this proposed definition of facilities-less competition, the difference between facilities-less competition and facilities-based competition is one of end-goal. The objective of *facilities-based* competition is to cause market entrants ultimately to build a complete parallel infrastructure to the one of incumbent carriers. In pursuing our analysis of facilities-less competition, we find that most implementations of successful competition in the marketplace were the result of forcing competition between players which significant disparities in the amount or number of facilities.

Perhaps one of the best examples of facilities-less competition is the successful implementation of long-distance telephony competition, where incumbent telephone carriers were required to aggregate traffic to a POI (point of interconnection) and modify their telephone switches to provide single-stage dialing equal access to a pre-subscribed long-distance carrier.

Facilities-less competition thus implies several notions:

- the concept of equal access,
- points of interconnection (POI) and
- back-haul or aggregation.

In implementing facilities-less competition, signals or traffic across a metropolitan area are aggregated by an incumbent carrier and is delivered to a designated POI. This results in a competition between an incumbent carrier and an entrant which has a **lesser** amount of facilities than the incumbent carrier. The POI provides equal access to incumbent subscribers and all parties are required to interface at POI before services can be delivered to the subscriber.

Finally, facilities-less competition does not require that entrants ultimately replicate the entire infrastructure of incumbent carriers, which is essential to see effective competition in the near future. The population is asking for competition now. Across the world, telecommunications policy makers have been given the mandate to implement competition as soon as possible.

In North America, where telecommunications policy is the most forward looking, we have yet to see competition develop in residential telephony despite several decisions requiring its implementation made more than three years ago. The current absence of residential competition is amplified by the fact that the focus of carriers, as regards residential service, has now completely shifted away from traditional circuit-switched telephony services towards high-speed Internet access.

Given this new focus, no new competitor will ever seek to build new facilities until such time as all opportunities to benefit from the existing infrastructure are no longer available. We argue that the outcome of existing telecommunications policy - which does not explicitly address the issue of high-speed Internet access in a uniform manner - prevents viable competition from emerging. This is so because the relevant form or forms of competition no longer consist of delivering identical forms of circuit-switched telephony, but rather in allowing for all the possibilities of the Internet to reach customers.

New Internet applications can replicate the services provided with circuit-switched telephony and conventional television equipment at a fraction of their cost. For example, the Session Initiation Protocol (SIP) Internet protocol is presently being used to unify Internet Telephony, Internet Television and Electronic Commerce applications across a common set of TCP/IP interfaces which are far more powerful than what the Signaling System 7 will ever perform.

For as long as it may be possible to interconnect with incumbent carriers to provide high-speed Internernet access services, the incentives to build new infrastructure will be ineffective.

We propose that the relevant agenda for regulators seeking to introduce competition in telecommunications is to apply the principles of 1) facilities-less competition, 2) equal access to packets and 3) minding the protocol stack to the issue of high-speed Internet access.

The reason for the importance of Internet access is that it now becomes the vehicle through which all new relevant IP – based services will occur. Circuit-switched telephony diminishes in importance as users of telephones, cable television and computers switch to newer technologies of transmission.

The rules for the provisioning of basic telephony service have been well established by the decisions regarding CLECs (competitive local exchange carriers). The new challenge is to accommodate the forthcoming migration of conventional circuit-switched telephony and cable television applications onto a high-speed Internet access infrastructure.

We invite the reader to follow *hypothetical* discussion between the president of a CLEC and its banker:

CLEC President:

My company will shortly be in the business of competing with a wire-based incumbent. As you may know, this market is well protected by law and the barrier to entry is very high. Although we expect to have some difficulty gaining access to electrical poles, our job will be facilitated by recent CRTC regulation forcing utilities to provide access to their infrastructure. Furthermore, we expect to benefit from market protection by CRTC regulation, such as contribution rules. Decision 97-8 does not require us to replicate all of the phone company's infrastructure because the CRTC has fixed the price of access to the phone company's infrastructure. We are currently into the third year of the five year term for unbundling of incumbent infrastructure. Since very little competition in the residential sector has emerged, we expect the CRTC to prolong the price fixing scheme for another 5 years, which will give us the chance to build the parallel infrastructure required by facilities-based competition regulation.

Banker:

As nobody makes money with local circuit-switched telephony services, explain to me how will your company provide high-speed Internet access ?

CLEC President

*The monopolies handed out by the Canadian government to telephone and cable television companies w*ere *for the delivery of basic telephone and cable television services. However, it is now clear that the existing telephone and cable television infrastructures can be used to provide high-speed Internet Access.*

As you are aware, it is now possible to provide basic telephony and cable television services through a high-speed Internet access connection. The fact that the incumbents are capable of providing high-speed Internet access services was never accounted for in current telecommunications policy. The rules of the game have changed and telecommunications has not followed. As you may be aware, recent decisions by the CRTC make it possible for competitors to provide high-speed Internet access services without building a parallel infrastructure. It is impossible for me to justify the construction of a parallel infrastructure until such time as I can no longer interconnect with the incumbent carriers. If my company builds a parallel infrastructure right away, we might have to face stronger competition from the ones who have used the existing infrastructure of the incumbent carriers. They could provide all services through a high-speed Internet access connection better than we could from building new infrastructure.

What we wanted to demonstrate by this hypothetical exchange between a CLEC and its banker was that, first, the possibility of delivering high-speed Internet access over the existing wireline infrastructure is a clear extension of the wireline duopoly beyond its planned purposes, and second, that telecommunications policy should not try to induce investment in needless infrastructure.

Additionally, the possibility of delivering telephone services over high-speed Internet access has seriously hampered the stability of the CLEC business model as introduced by the CRTC in Decision 97-8. Finally, recent telecommunications regulations have worsened the problem as it now appears to be possible for competitors to provide high-speed Internet access over the infrastructure of incumbent telephone and cable television companies. The incentives to build a parallel infrastructure have vanished, until it becomes clear how access to the high-speed Internet access facilities of incumbent carriers will be governed.

The concepts of facilities-less competition, equal access to packets and minding the protocol stack were originally developed for the purpose of finding a common way of thinking about third-party access to high-speed facilities of cable television carriers within existing telecommunications policy.

Unlike the telephone network, cable television distribution networks are based on a shared infrastructure which make physical interconnections more difficult. Cable television carriers maintain that interconnections at a higher layer are required if they are to meet common carrier obligations. Because incumbents require that intereconnections in a cable environment be made at a higher layer than the physical layer, potential competitors can have no control over physical facilities. For this reason, facilities-less competition, on a permamanent basis, is thus absolutely necessary.

We conclude that in all current occurrences of successful introduction of competition, the mechanisms[12] of facilities-less competition have been employed. However, facilities-less competition has been considered as a short-term incentive to create facilities-based competition. We urge this policy be reconsidered in the light of new facts.

[12] Examples: Source-based routing, various tunnels, MPLS, RED, WFQ, DOCSIS 1.1, RSVP, etc.

Public infrastructure versus private infrastructure

The choices before us are clear. On the one hand, we can institute an interconnection regime where private ownership of carriage facilities continues. In such a regime, a high level of scrutiny is maintained over interconnections, so that incumbents do not privilege their signals at the expense of others. Regulatory scrutiny would apply to all layers of the protocol stack and to the physical interconnections. Equal access and equal treatment of packets would have to apply to all aspects of such a system. Ownership of the physical layer would not imply a right to discriminate against other uses at higher layers.

The alternative is to lay down a public infrastructure of optical fiber, akin to roads, managed by municipalities, school boards, provinces and the federal government, as appropriate. Service providers could light up portions of the bandwidth, and pay rent to the government for the use of the fiber. This path might be adopted if regulators and politicians conclude that Canada will not achieve sufficient Net connectivity and the benefits of the open-architecture Internet model fast enough under an interconnection regime like the one being discussed in this paper. The choice of one or another path, or the combination to be followed, are matters of practical wisdom.

2. Equal Access to Packets

Equal access to packets is a logical extension of the concept of facilities-less competition which introduces points of interconnections (POI) to segregate traffic between the infrastructure of an incumbent carrier and its competitors. It follows that the concept of non-discriminatory interconnections requires that the POI cannot be used by the incumbent carrier to exercise self-preference.

In an Internet environment, the network uses packets to transmit signals. Packets are groups of hundreds of binary digits which are used to represent portions of complete messages. Each packet is addressed to a destination and computer hosts at each end-point decide when and whether or not to retransmit a packet if it is dropped by a particular Internet router.

In the telephone environment, the notion of equal access was first introduced when telephone companies were forced to modify the software in their telephone switches to implement equal access for long-distance carriers. It followed that local exchange carriers back-hauled their traffic to a single point of interconnection in a metropolitan area and used the point of interconnection to provide their own services.

In Canada recently, one particular cable carrier was building its network in such a way that the point of interconnection was only going to be used by competitors. In addition, completely different routing policies would by applied to packets of the incumbent cable carriers than to the packets of the various competitors. In this situation, the point of interconnection was built as a firewall and not to provide equal access.

Equal access to packets is a corollary to the idea of facilities-less competition. We argue that incumbent carriers should be required to use the same point of interconnection as competitors, and that by logical extension the same treatment must be applied to packets of incumbent carriers as to competitors' packets.

We further propose that the means of control of the equipment which manipulates packets must not be in the sole control of the incumbent carrier.

Finally, we also propose that the choice of the technology and the equipment which implements this technology is not at the sole discretion of the incumbent carrier and must be subject of industry-wide agreements. In requiring incumbents to use and share with competitors all points of interconnection, to design them cost-effectively, and to ensure that the treatment of packets is similar across all parties involved, facilities-less competition can be the source of effective competition.

In an Internet environment, unless the concept of *equal access to packets* is formally introduced in all new decisions, policies and regulations, and policed in the manner proposed above, the complexity of the mechanisms[13] involved will permit incumbents to implement self-preference, and will completely surpass the ability of any regulator to police the technology and the architecture of points of interconnection.

3. Minding the protocol stack

We propose that regulators must now concern themselves with minding the protocol stack. The principle of equal access to packets requires that regulators and the industry must be concerned with interconnection through software, which occurs at a higher layer than the the physical apparatus of connection and transmission.

Current telecommunications policy is concerned with interconnection of facilities at the *physical layer*. The opportunities for self-preference in the interconnection at a higher layers are inherent feature of a software-driven system.

In the particular case of third party access to high-speed facilities of incumbent carriers, the interconnection between Internet Service Providers and Cable Carriers is based on a common language called TCP/IP. However, TCP/IP is not the native language of the physical equipment which is used by either the ISP or by the cable carrier. In the case of the cable carrier, the cable modems speak a language, or a physical layer protocol, or media access control protocol called DOCSIS, which stands for Data

[13]Examples: RED, WFQ, GRE Tunnels, MPLS, RSVP, DOCSIS 1.1, Diffserv, Intserv

Over Cable Service Interface Specification.

DOCSIS is a variant of the Ethernet media access control protocol used over category five "blue" wires in all of our offices. Cable television carriers are incapable of providing interconnections to their facilities using DOCSIS since this technology was never designed to be shared amongst multiple service providers. In this sense, the interconnection between ISP's and cable carriers cannot be called third party access or even less third party Internet access, because no connection to the Internet is provided in this environment. Rather, such an interconnection should be called what it is: a TCP/IP version 4 network layer interconnection to DOCSIS networks of cable carriers.

Incumbent carriers can be granted an unfair advantage when they are allowed to meet their common carrier obligations by forcing a certain type of network-layer interconnection onto competitors. It is through the capacities built into the software protocols – or not built in – that the design philosophy of the Internet can be achieved or frustrated. We have argued that the great economic boom that has accompanied the Internet is based on a design philosophy of end-to-end connections, which mean that the system is radically underspecified from an engineering point of view as to what goals it should accomplish. This has allowed tremendous innovation in products and services, because no one has controlled what the system would be used for. Every innovator could bring his product and service to market.

The limitations or design of the network layer interconnections have an enormous impact on what gets designed, produced and sold in the new era. Redefining network interfaces redefines markets and services, redefines winners and losers, abolishes or reinforces monopolies. It is that important.

For example, in the Internet over PSTN environment, ISPs based a significant part of their business model on their ability to redefine the interfaces of the PSTN from usage-sensitive circuit-switched dialup voice interfaces into usage-insensitive permanent V.90 packet-based Internet Protocol version 4 over point to point protocol (PPP) interfaces. In the context where the public network of incumbent carriers now becomes based on the Internet Protocol version 4 (IPv4), it is very likely that interfaces with competitors will need to be redefined again, and that incumbent carriers will find that IPv4 cannibalizes their vertically integrated service offerings. For example, Internet Service Providers will most certainly seek to redefine the public IPv4 Internet interfaces into Internet Protocol version 6 (IPv6) at the request of users fed up with not being able to get enough IP addresses for all devices on their home area networks.

In the case of cable television carriers, it is well known that their common carrier obligations could by met by installing the equivalent of main distribution frames into their head-ends in order to share portions of the cable spectrum with competitors. With such equipment, the equivalent of CLECs in the cable environment would be free to specify the nature of the technology being used.

However, current decisions regarding high-speed access concern themselves only with prices, not with the nature of the interconnection that is provided. Tariffs are a means for cable carriers to exercise tighter control over their facilities than otherwise could be possible with physical interconnections. Furthermore, with network layer interconnections, the incumbent carriers are granted the opportunity to segment their own shared network and to control the ratio of asymmetry between the forward and the downward links on the infrastructure. In doing so, the incumbent carrier still possesses several opportunities to protect the cannibalization of existing services.

Incumbent carriers are relieved from having to meet common carrier obligations to competitors by only having to provide network layer interconnections. Although an effective policy of competition would see them allocating spectrum over their wireline infrastructures, on the principle that they should deliver the bits, not control the definitions of services, in the absence of such a policy, network layer interconnections are subject to public scrutiny.

Conclusion

The economic boom unleashed by the Internet relies on a particular idea embedded in it: the lack of central planning of what the network should do. The previous pre-computer concept of communication networks permitted only the owner of the network to design the services that could be generated from the system, and the result was decades of technologial stagnation in telecommunications, relative to what computers have enabled. The two systems are in collision, and one must prevail. Apparently abstract things like network layer interconnections become the means whereby one system prevails over another, or does not. Accordingly, government's traditional role of constraining monopoly power in telecommunications takes on fresh and urgent relevance.

We conclude by saying that telecommunications policy should now be expressed in terms of facilities-less competition, equal access to packets and minding the protocol stack. Failure to follow these precepts in an Internet era would be fatal to competition. The opportunities to exercise self-preference in an Internet environment are multiplied by the complexity of protocols, the rapidity of software-driven change, and the lack of understanding by almost anyone of these issues. Without adherence to these precepts, incumbents may surpass the ability for telecommunications policy makers to ensure viable competition.