

WRITTEN TESTIMONY OF CHRISTOPHER S. YOO

**Professor of Law and Communication
Founding Director, Center for Technology, Innovation, and Competition
University of Pennsylvania**

Hearing on “Net Neutrality and Free Speech on the Internet”

Before the Committee on the Judiciary of the United States House of Representatives

March 11, 2008

Mr. Chairman and Members of the Commission, I am grateful for the opportunity to testify here today. As a Professor of Law and Communication and the Founding Director of the Center for Technology, Innovation, and Competition at the University of Pennsylvania, I have been studying these issues for the better part of the past decade. I would like to make two basic points to day.

The first is that the recent surge in Internet traffic caused by new video technologies, such as YouTube, has the potential to overwhelm network providers’ ability to add capacity to meet that demand.

The second is that much of the modern Internet is increasingly dominated not by person-to-person communications, but rather by mass media content. There is a long line of Supreme Court precedent acknowledging the important free speech benefits of preserving editorial discretion when media content is involved, both to help sort through the avalanche of content available on the Internet and to screen out bad speech, such as spam and pornography.

The Internet has changed dramatically since its origins as a way for university-based scholars to communicate with one another. Since the Internet backbone was privatized in the mid-1990s, the number of Internet users has exploded. This dramatic increase in the number of users has been accompanied with a similarly dramatic increase in the diversity of ways people

are using the Internet. The relatively simple applications that dominated the early Internet, particularly e-mail and web browsing, are being superseded by a host of more complex and demanding applications, such as Internet telephony, streaming video, online gaming, and virtual worlds. Not only do these new applications demand more bandwidth; many are also much more sensitive to delay than were e-mail and downloading webpages.

The evolution of the Internet has made networks increasingly difficult to manage. The problem is exacerbated still further by the uncertainty about how quickly Internet traffic will likely grow in the future. Between 1996 and 2002, Internet traffic doubled every year. Since that time, the rate of traffic growth has stabilized at roughly 50%-60% per year. Thus, even the most conservative estimates project that network providers will have to undertake massive new investments just to keep pace with traffic growth. Some experts have predicted that the rise of new Internet video technologies like YouTube may once again cause Internet traffic to double each year. If so, Internet video will unleash an “exaflood” that will slow the Internet to a crawl by 2010 unless an additional \$137 billion is invested in network capacity.

The reason is that the Internet is subject to congestion. In other words, the speed depends not only on how much capacity any particular user is using, but also on how many other people are on the system at the same time. There are typically two solutions to the problems of congestion. Network providers can build excess capacity to make sure that there is enough bandwidth to ensure against delays when traffic levels peak. The problem is that adding capacity is a solution can be quite expensive, can slow rural deployments, and may ultimately be self defeating. The alternative is to allow network providers to engage in network management. The fact that predictions are never perfect and that capacity cannot be added instantaneously necessarily means that network management will always remain a valuable means for preserving

the value of the Internet. Foreclosing that alternative would place network owners in the impossible position of having to be perfectly prescient in their predictions of exactly how much Internet traffic will grow and exactly where geographically extra capacity will be needed.

The choice between expanding capacity and network management can be illustrated through a more familiar resource that is also subject to congestion, which is our Nation's road system. Like the Internet, traffic on the road system tends to fluctuate. Thus, even though the posted speed limit may be 55 or 75 miles per hour, the speed particular drivers will actually achieve depends not only the total number of miles they drive on the number of other drivers on the road at any particular time. In addition, like the Internet, congestion on the road system is not uniform. Although most portions of the road system will be uncongested, there will be the occasional hot spots.

How do transportation engineers alleviate congestion during times of peak traffic? One solution is to add more lanes. This solution suffers from a number of drawbacks, however. As an initial matter, maintaining excess capacity makes building roads considerably more expensive. In a world in which the network resources are privately provided, the inevitable consequence is to increase the number of customers needed for any network expansion to breakeven, which in turn inevitably limits the ability to expand into less populous areas. Moreover, traffic engineers' predictions of exactly where and how many additional lanes will be needed are never perfect. Absent traffic management, any errors due to underestimation would take years to correct.

Even more importantly, the evidence suggests that expanding capacity may ultimately prove futile. The addition of lanes simply stimulates development at the ends of the road until

the new lanes become congested as well. In short, demand will expand to fill all available capacity no matter how much is added.

The alternative approach to managing peak traffic is engage in some type of traffic management. Classic examples include putting stoplights at Interstate on-ramps that restrict the traffic entering the system, giving ambulances and other traffic with a higher social value priority over other traffic, providing drivers with fast-pass options at toll plazas, and reserving lanes for buses and high occupancy vehicles. Each of these approaches represents a form of nonneutrality. And yet, each is regarded as uncontroversial.

I do not mean to push the analogy between the road system and the Internet too far. There are some critical differences between them. For example, unlike traffic on the road system, in which congestion rises and falls rather slowly and steadily over time, traffic on the Internet is extremely “bursty,” in that long periods of inactivity are punctuated with extremely brief but intense periods of heavy bandwidth usage. This makes the network management problem considerably more complex and calls for different tools.

In addition, the Internet is comprised of radically different technologies which vary widely in their susceptibility to congestion. Telephone-based technologies allow telephone companies to give end users dedicated connections all the way to the central facility they maintain in each city known as the central office. As a result, telephone-based technologies are not subject to congestion at the neighborhood level. In contrast, cable-based technologies by their nature do not offer such dedicated connections. Instead, end users share their connection with roughly 300 other users located in the same neighborhood. This makes cable modem systems much more susceptible to neighborhood congestion than telephone-based systems. Unsurprisingly, cable-based systems must take different steps to manage their networks. These

problems are even more acute for wireless broadband. Not only do wireless users share bandwidth in much the same way that cable-modem users do. The bandwidth available to wireless technologies is much more limited than that available to any other technology. As a result, wireless providers frequently give higher priority to telephone calls and other applications that are time sensitive while holding e-mail and other traffic for which delays of a fraction of a second are essentially unnoticeable.

Perhaps the most important difference between the road system and the Internet is the presence of bandwidth hogs. In the road system, each driver causes roughly the same amount of congestion. On the Internet, the situation is quite different. Network providers estimate that as few as 5% of end users represent between 50% and 80% of the network's total usage, and many applications are designed to increase their usage as long as capacity is available. The question in such a world is not whether congestion will occur. The existence of applications that increase their usage to fill all available bandwidth makes that inevitable. The question is whether the costs of those congestion will be borne by all users or only by the handful of users responsible for that congestion. Under these circumstances, requiring those most responsible for congestion to bear a greater percentage of the costs would be both good network management and fair from a consumer standpoint. Any other system would in effect require low bandwidth users to cross subsidize the network usage of a handful of bandwidth hogs. Charging differential rates is thus more properly regarded as a way to keep fees down for end users who simply want to maintain websites.

It is for this reason that every panelist that testified at the FCC's February 26 hearing on network management agreed that some degree of network management is inevitable. The problem is that the reasonableness of any particular approach to network management varies

from technology to technology and even within any particular technology varies depending on the amount of traffic in the network at any particular location or at any particular time of day. The problem is complicated still further by the fact that the technology underlying the Internet is undergoing constant and rapid change. Any categorical solution that might be mandated today would all-too-frequently be overtaken and rendered obsolete by technological change. In the worst case, a mandate might retard new technological solutions from emerging in the first place.

My second larger point is that the nature of the Internet is changing in ways that the current debate does not yet take into account. Proponents of network neutrality repeatedly cite telephone-based precedents, such as the *Computer Inquiries* and *Carterfone*, which were developed to govern person-to-person communications. In so doing, they overlook the fact that the modern Internet is not dominated by person-to-person, but rather has become an important source of mass-media content. As a result, end users have turned to a wide array of intermediaries, such as search engines, bloggers, and update services, to help them sift through the avalanche of desirable Internet content that grows with every passing day. They also depend on intermediaries to protect them from undesirable content, such as spam, viruses, and pornography.

The question is thus not whether some network provider will serve as an intermediary. The flood of Internet content makes that inevitable. Instead, the question is who will play that role. Whoever plays that role must walk a careful line, in that they will be asked to screen out only that content that is truly undesirable while allowing all other content through. Consider, for example, Verizon's statement at the FCC's February 26 hearing. Verizon revealed that it has declined to activate short code campaigns that were designed to distribute wallpaper with nude images or ringtones containing profanity or racial slurs, would be used solely for spam, or would

charge inordinate fees. Any legal intervention must thus be carefully tailored not to prevent exercises of editorial discretion that benefit consumers, keeping in mind the Supreme Court's famous admonition in *Cohen v. California* that "one man's vulgarity is another's lyric."

The likely difficulty is illustrated by the positions taken by the Christian Coalition. On the one hand, the Christian Coalition has endorsed network neutrality. On the other hand, it has asked the FCC to preserve network providers' ability to screen out child pornography and obscenity and to refer any instances that they may discover to law enforcement authorities.

The result is that we are asking network providers to strike a careful balance, both in terms of building the Internet of the future and in terms of managing the ever-growing deluge of content that is available online. Prohibiting network owners' ability to play this role would simply shift the emphasis to search engines and other intermediaries, which no doubt would simply cause the debate to shift to "search neutrality." In this sense, the history of Google's success is quite revealing. Google displaced existing search engines like AltaVista by coming up with an algorithm that prioritized search results in a different way. This raises serious questions over whether any particular algorithm can ever be considered truly neutral and suggests how network providers can compete in the way they exercise editorial discretion.

It is for this reason that I have long opposed imposing categorical, before-the-fact prohibitions on any particular conduct. Many forms of network management are beneficial, and categorically precluding any particular form of network management threatens to choke off many of the innovative solutions made possible by new technological developments.

At the same time, the economic literature reveals the existence of limited circumstances in which network owners might use their economic position to disadvantage consumers. Although such circumstances would be relatively rare, such harm remains theoretically possible.

It is for this reason that I have long advocated taking an after-the-fact, case-by-case approach to regulating network neutrality. Such a position would have strong support from Supreme Court precedent. It would also strike the proper balance between protecting consumers and giving dynamic industries like the Internet the breathing room needed to respond to changes in the underlying economics and technology. Although it is conceivable that antitrust courts might play this role, I believe that they are institutionally poorly suited to the type of ongoing supervision that overseeing an access mandate would require. It is for this reason that I have argued that a regulatory agency like the FCC is in a better position to oversee such an analysis. The Supreme Court's *Brand X* and *Trinko* decisions both agree that the FCC has both the authority to oversee such a mandate and is in the best institutional position to do so.

I am heartened by the fact that an ever-growing number of industry participants and scholars have endorsed the approach I first proposed. The vigor with which the FCC has pursued allegations of improper network management suggests that the regulatory structure may already be in place to ensure that consumers are both protected and able to enjoy the Internet's tremendous promise in the future.