



Network Neutrality: Words of Power and 800-Pound Gorillas

DAVID D. CLARK

Massachusetts Institute of Technology

The debate over network neutrality means different things to different stakeholders. To advocates for the Internet, it has to do with preservation of the Internet's open nature. On the other hand, to large service providers, it has to do with the viability of revenue streams, and in particular the future of television. The current structure of the television industry is much more vertically integrated and "closed" than the Internet, and this structure has sustained a well-understood value chain with stable opportunities to profit from content delivery, advertising, and so on. In contrast, the open design of the Internet limits the ability to profit from the delivery of content, and as the Internet more and more becomes the platform to deliver video, the threat to revenues is obvious. It is this collision between different revenue models, each valid in its own context, that provides the drive for the debate over network neutrality. The history of Voice over IP (VoIP) is a cautionary tale about the disruption of existing revenue models by the Internet, and those who count on the current revenue models for television are preemptively moving to defend themselves. Any attempt to intervene in this battle must first take a realistic view of what the options are for the stakeholders and how the battle might play out, and must not adopt an over-simple or over-constraining view of what it means to be "open."

The Power of "Open"

The one-word bumper-sticker that captures the character of the Internet is "open." The Internet is open in many ways: it is an open platform (at the Internet layer) with low barriers to innovation. The standards are open in that they are free of IP or licensing limitations. The protocols allow "anyone" to become an Internet Service Provider (ISP) and join the club. The emergence of these features is not an accident: the designers of the Internet valued and defended them. And many connections have been claimed between these features and the broader virtues of "open," which is a very powerful and positive word that can be inserted as an adjective in front of many nouns to express a positive value: open borders, open conversation, open countenance, open access, open markets. "Open" is good, "closed" is bad. The Internet, to a large extent by design, but also by happy accident, has a technical character that

David D. Clark: ddc@csail.mit.edu

Date submitted: 2007-06-07

Copyright © 2007 (David D. Clark). Licensed under the Creative Commons Attribution Non-commercial No Derivatives (by-nc-nd). Available at <http://ijoc.org>.

aligns it with this larger set of values that equates, in the extreme, with the power of democracy and freedom.

Many advocates of network neutrality are fighting to defend openness. But not everyone in the debate is fighting over the social value of the Internet. The 800-pound gorillas, the large industrial players, are fighting over revenue models and the future of an industry. And right now, what they are fighting over is the future of television.

It is more or less accepted today that Internet technology will be the platform to deliver entertainment video to the consumer, whether the delivery is over "the open Internet" or over private, dedicated facilities based on Internet technology. Already, the Internet is providing the consumer with a variety of video opportunities including YouTube, official and pirated television shows, pirated movies, traffic cameras, day-care and nanny monitors, and almost any imaginable sort of webcam.

While video on the Internet is not yet challenging the traditional entertainment product delivered over the air or over cable, television and the Internet are on a collision course. And the future will indeed be a collision. The cable systems are not "open" in the same ways that the Internet is. While spectrum limitations have driven the regulation of over-the-air broadcast, and there are some obligations placed on cable providers to carry local content, it is largely true that the choices of content available to the consumer are the choices determined by the owner of the distribution facility. We have accepted the "somewhat closed" nature of the TV distribution channel, and accepted as normal the resulting industry structure and revenue models. We accept that the cable provider picks the content to distribute, buys it "wholesale," marks it up and sells it at retail.

And it is the collision of these two models that is one of the major battles lurking inside network neutrality. What is the nightmare of Verizon, for example, as they invest heavily in their fiber to the home product, FIOS? FIOS allows Verizon to offer a TV product that can compete head-to-head with the cable providers. So one might think that their nightmare is that they cannot capture market share. No, the nightmare is much worse than that. It is that just when they are in a position to compete for market share, the market itself vanishes, to be replaced by "free" content, or by content provided by third parties which access providers cannot mark up and sell at a profit, but which they are expected to carry for "free" over their "open" infrastructure.

I have put these words in quotes to stress the cartoon-like nature of the story I have told. We know that even if the Internet is open, it is not free. We know that the Internet is not fully open today, in lots of ways. But the fear is nonetheless real, and nonetheless valid. The cable industry today purchases content at wholesale and sells it at retail. There is no equivalent business model in the Internet, whether the content is music, video, or the written word. The value chain for Internet video is very complex, with many sorts of content creators and distributors, but the Internet Service Provider that provides access to the consumer is just expected to "carry the bits."

Looking Back at Earlier Collisions

The television industry is not the first sector to collide with the Internet. The recent past offers obvious examples from which some useful insights can be drawn. Internet telephony, now called VoIP, is illustrative of what happens when pre-existing industry sectors collide with the Internet. It was clear around 1980 that the Internet could carry voice streams. It was clear by the early 1990s that the Internet could carry voice “at scale”—it could have enough capacity to carry as many phone calls as the existing phone system, but there was still a lot of doubt then as to whether it could be reliable enough, offer a quality service, etc. Phone companies, at that point, were not too worried about either the technology or the threat to their business models.

VoIP first emerged as a disruption in the long distance market, and I would argue that it was first seen not as a disruptive technology but a means to challenge the regulatory framework of the day, which still had substantial access charges flowing from the long distance provider to the local access provider. But even as the incumbent long distance providers made statements about conversion to IP technology (and thus avoiding access charges), they may not have grasped just how dramatic and traumatic the transformation of their business model might be, with “free” voice calling eroding their traditional per-minute billing. The Internet has disrupted the traditional billing model of the phone companies, and contributed to a decline in demand for residential access lines. And just as the Internet contributed to the decline of the revenue model of the telephone company, it might destroy the revenue model of those that provide access to television content.

The music industry has undergone a similar disruption as it collided with the Internet, and in particular with the music sharing applications such as Napster. In this case, the issue did not involve access circuits but intellectual property rights (IPR). The music industry may have been slow to respond to the threat/opportunity of the Internet, and seems to be fighting a rear-guard action to protect their intellectual property, as well as the viability of their existing distribution channels.

VoIP and music stand as cautionary tales to those who look into the future and see that the Internet might invade their territory. If some stakeholders might have been caught napping by the collision of the Internet and telephony or music, this will not happen with television. The battle lines are being drawn early and often. And the issues with television are similar, in general terms, to those of telephony and music. Television is a product where the “service,” not just the access, is the basis of the revenue. It raises issues of intellectual property, and fears of the providers in the current production and distribution channels that the current industry structure (and its revenues) may somehow fall apart when confronted with the open nature of the Internet.

So the debate about network neutrality is a proxy for a debate about the collision of different models, all valid in their own sectors — the open nature of the Internet, and the tied content-conduit model of the entertainment industry.

There are other coming battles, related to the future of TV but to some extent distinguishable from it. Another one is the future of advertising. The cable industry has a tradition of profiting from

advertising by doing local ad insertion. Right now, on the Internet, the business (and profits) of ad insertion belong to companies that have nothing to do with network facilities, and Internet Service Providers have no opportunity to participate in this revenue stream.

The collision described above is inevitable and tectonic. The industry players are large, skilled and well-funded. They are not going to give up their business models or their profits without a fight. And while advocates may be unwilling to articulate an outcome that is anything less than total victory for their side, we must accept, if we can be dispassionate observers, that neither side will emerge with its virtues or its profits intact. We should thus be willing to contemplate what the shape of the outcome might be, what the factors are that will shape the outcome, and what specific considerations we most care about. With this in mind, I offer some thoughts about the present and the future.

Fear of Gorillas

One issue that lurks under the covers of network neutrality is fear of market power. We can now see the emerging industry structure of broadband access providers, which in the United States is often a duopoly between a “telephone” and a “cable” provider, and in other parts of the world may be an effective monopoly. Competition can be found, of course, but there are fears about the possible market power of these access providers, and whether regulatory intervention will be appropriate to limit that power. (I am not saying that the fact of market power is proven. I am saying that the fear itself is real and present, and needs to be accepted as one of the drivers of today’s actions.)

While debates about network neutrality are now breaking out around the world, the term (and the tension) seems to have arisen in the U.S., and one could speculate that the reason for this is that we have abandoned the idea of increasing competition through facilities unbundling, and see (to some extent) the outcome of that decision, whereas other parts of the world are following the path of encouraging competition at the retail level through unbundling, and are thus hoping that the issues of market power at the retail level will be less pronounced.

As a counterpoint, it might turn out that we have misunderstood where market power might emerge. A firm like Google might just establish such a hold on Internet search, Internet advertising, Internet e-mail and other content that they come to dominate their sector and set the terms of the collision. There is fear on all sides.

Avoid Collateral Damage

The debate over television (and older debates about telephony and music) could well spill over into other aspects of the Internet. An important goal is to avoid collateral damage from the battles. One obvious concern is the future of the Web. The Web is a wonderful example of the power of “open”— it combines the technical and social forms of openness in a way that has transformed society. Anyone can make a Web page, and while content providers may choose to limit (or charge for) access to their content, no third party can block (or profit from) a user’s access to a Web page, except in the case of conservative governments that have used their power to “close” the Web. It would be a disaster for the Web and for

society if this open nature were eroded by a fight over television. But that might happen. One of the so-called virtues of the Internet is that “bits is bits,” and all bits are equal. So even if the defenders of the Web would agree to throw entertainment content to the wolves to protect the open nature of their own space, they are caught by the consequences of convergence (another wonderful buzzword) — how can anyone draw a line between what the Web does and what television does when they are all “just bits,” and when the boundary between the two blurs and vanishes when you look closely. The Web can deliver video content—where is the line? It is not there.

Being Realistic About Traffic Discrimination

If access providers like cable and telephone are going to profit from content that is distributed across their network, one approach they might take is to detect different content as it is being sent and either bill for it, or discriminate against it by blocking it or degrading it. So the open nature of the current Internet has been equated to a lack of discrimination among different sorts of traffic, and the advocates of network neutrality often declare their policy as “no discrimination”. The word “discrimination,” of course, is a loaded word, although to an economist it is not a dirty word, but a term that describes legitimate business practices.

There are two points about this perspective on network neutrality. One is that discrimination in service quality among traffic flows is not the only way an ISP could arrange to profit from content distribution, and the other is that there are already many forms of accepted discrimination going on in the Internet today. We need to explore these two points as a reality-check on the “no discrimination in service” characterization of network neutrality.

An obvious form of current discrimination is between traffic that needs to be delivered by a strict deadline (so-called “real-time” traffic, such as VoIP), and traffic that just needs to get there as soon as (and as fast as) possible (so-called “best-effort” traffic). Some ISPs distinguish between these two sorts of traffic today, to improve the quality of the real-time traffic. (Preferred treatment might be offered to the real-time traffic of services that *they* offer, but not the real-time traffic of third-party providers, which is an example of how service discrimination can translate into business discrimination.) The mechanisms used to implement this sort of discrimination are usually called quality of service (QoS) mechanisms. As another example of current discrimination, some Internet providers block various applications, often based on an argument about security, and degrade or limit various applications, such as peer-to-peer music sharing programs.

As a technical mechanism, QoS seems to be beneficial. It directly addresses the real performance requirements of different sorts of Internet traffic. By allowing different sorts of traffic to ride on the same network, implementation of QoS might side-step an even more dramatic form of discrimination, which is to build multiple, parallel versions of the Internet, each supporting different service qualities, and separately selling access to each. We already see this sort of behavior today in some advanced broadband access networks, where the ISP uses Internet protocols and technology to implement two different versions of the service over the same access connection: one a connection to the open Internet supporting only best-effort service, and the other a closed network similarly based on Internet protocols,

supporting QoS, but only available to the provider of the access network. This second network is then used to deliver services such as VoIP or video that require service quality commitments. Is this form of discrimination more or less tolerable than variable access to different QoS options on a single, public Internet?

This reality begs the question of whether we can find a set of rules that might distinguish between “good” or “acceptable” forms of discrimination, and “bad” discrimination. Unless we can find a bright line, using regulation of discrimination to define acceptable behavior may cause more trouble than it cures.

Interconnection

The old joke about gorillas goes as follows: “Where do 800-pound gorillas sit?” “Anywhere they want.” But where do 800-pound gorillas fight? They fight where they have contact with each other, or (in Internet terms), where they interconnect. ISPs connect with each other, and they interconnect with other sorts of entities, such as content providers. The debate about network neutrality is often portrayed as protection of the consumer experience, and regulations are often framed in these terms. But it should be clear that what happens to the consumer is a byproduct of how the battles among the large stakeholders are resolved. So there are two ways to contemplate the problem of neutrality — describe our vision of how to protect the consumer experience and use that to limit how the gorillas fight, or to address directly the nature of interconnection.

To start, let us look at the shape of interconnection among ISPs. ISP interconnection is often portrayed as very simple — either a small ISP buys so-called transit service from a larger ISP (which gives that smaller ISP access to all of the Internet via the larger one), or two ISPs agree to “peer,” with no payment between them, in which case they connect with each other so that each can deliver to the other the traffic that is destined for it. But this over-simple story masks a wide range of business practice. Different ISPs will agree (or not) to peer with different of their potential inter-connection partners. ISPs may offer reduced costs for transit based on other business conditions that apply. Some ISPs may try to charge for the privilege of peering. So there is a great deal of discrimination that can be found in today’s practice of interconnection.

Given that reality, how do ISPs connect to other sorts of actors, such as providers of content (including video)? Large content providers generally do not deliver content into the Internet from one point—one gigantic server located somewhere. They tend to use multiple sources, distributed around the Internet and attached near the consumers, and perhaps their own distribution facilities. This reality can lead to two very different conceptions of the outcome. In one outcome, each of these sources is attached near a set of consumers (e.g., directly on an access network), and pays for this connection. In this model, money flows from the content provider to the various access providers. In the other outcome, the content provider has his own distribution facilities, and looks like just another network, with links and multiple interconnection points. When this network connects to other ISPs, should the expectation of revenue-neutral peering apply? What is the difference between a network and a set of content sources? Both are connections over which traffic flows into an access network. This situation suggests that the simple model

of revenue-neutral peering is not stable, and brings the discussion of network neutrality away from the treatment of the consumer, and toward a consideration of interconnection discrimination, where ISPs may try to force content providers to pay to deliver their content into the network, and content providers may try to demand revenue-neutral peering.

Consumer Pricing Makes the Problem Worse

The concerns of the access providers today about cost recovery and third party video are aggravated by a factor that itself has nothing to do with network neutrality, but which adds insult to injury.

Today, the dominant pricing model for consumer broadband access is a flat monthly rate, dependent on the peak access rate but not on the total amount of traffic sent. Given this pricing model, not only is there no way for the access provider to profit from the delivery of video, he cannot even charge more to deliver the additional bits. Since the service is flat rate, if the consumers start watching video, the access providers' costs go up as he is forced to re-engineer his network for the increased traffic inside the network, but he sees no new revenues. Either the pricing models will have to shift away from flat rate independent of usage, or some other party (e.g., the content provider) is going to end up carrying some of these costs.

A Different Example of Content Discrimination

Since the discussions about network neutrality have sometimes equated neutrality with lack of service discrimination, it might be helpful to have an example of a form of discrimination that has only to do with price, and not with service quality. Imagine that over time, the access ISPs shift their customers away from flat rate pricing for unlimited usage to a tier of rates associated with different monthly usage caps. (Monthly usage caps are already emerging in the broadband cellular market, and there is no reason to believe that they will not emerge in the wireline market.) Under this circumstance, a user that wants to download a large amount of video will have to purchase a large monthly quota, and this cost might discourage the emergence of Internet video. To mitigate this situation, the access ISP might go to certain content providers, and negotiate a deal by which the provider pays a higher fee to deliver the content into the network, and the access ISP delivers the video to the consumer without deducting the material from the monthly quota. With this set of relationships in place, the consumer will perceive that there are two sorts of video, the kinds that do and do not use up his monthly quota. Now the ISP can use any power he has in the market to negotiate different contracts with different providers, or indeed to refuse to offer this option to certain, less-favored providers. This situation is indeed discrimination, but not about service quality. If this behavior were deemed unacceptable (and I do not take a position on that question), then regulation that is put in place to cover unacceptable discrimination should be general enough to cover this sort of behavior, as well as service discrimination.

Finding the Middle of the Road

I believe that there are examples, perhaps hypothetical at this point, of “non-neutral” behavior so abusive that they should not be tolerated. Given the concern about real market power, it seems unlikely that we can totally ignore the possibility of behavior that materially distorts the valuable attributes of the Internet. The challenge to those who would defend network neutrality is to recognize the complexity of the current situation, and find dividing lines —“bright lines”— that separate what is acceptable and unacceptable. Especially if these limits are to be built into *ex ante* laws, they have to be clear and workable. Perhaps an approach that depends on *ex post* administrative review will allow more flexibility and the emergence of “neutrality case law,” but even there we should expect calls for clarity and regulatory predictability from all sides.

At the moment, the ball is in the court of those who favor network neutrality. They have the responsibility to propose a set of definitions for the dividing line between acceptable and unacceptable discrimination. The issue is complex, and Washington likes simple cartoons. So how will the debate make progress? Most of what we have seen so far (in my opinion) either greatly overreaches, or is so vague as to be nothing but a lawyer’s employment act. Those of us who may agree that some form of bounds on bad behavior would be important have acquired the job of defining the bound.