

Inside Risks

Routing Money, Not Packets

Revisiting network neutrality.

THERE IS AN old conversation between two friends, a very prominent economist and Dave Clark (one of the foundational architects of the Internet) that is often recounted by Clark:^a

Economist: “The Internet is about routing money. Routing packets is a side effect. You guys really screwed up with the money routing protocols.”

Clark: “But we didn’t design any money routing protocols!”

Economist: “That’s what I said.”

While this was lighthearted banter between friends, the economist had it right 15 years ago—the crux of the issue behind what is broadly termed as network neutrality was captured in the conversation. The Internet, which began its life as a network offering communication services between primarily non-profit entities such as academic and research institutions has now evolved into the backbone of unimaginable amounts of commerce.

While it is tempting to think of every packet being equal, the revenue each packet generates is different.

The raging debate over network neutrality can be explained by this changing nature of the Internet, the players involved and the economic motivations. The FCC recently proposed Open Internet (the official FCC term for network neutrality) regulations, and invited public comments on it. Four million people responded, mostly in favor of the regulations. Finally, on February 26 the FCC had a vote (3-2) in favor of those regulations. While there was much celebration and joy among the activists on the success of the vote,

there are some technical reasons we need to consider. In this column, I lay out those concerns.

The Problems with the Ruling

The FCC proposed three “bright-line rules,” namely “no blocking, no throttling, and no paid prioritization.” One of the prominent reasons behind those rules is the very public dispute that Netflix had with major ISPs. However, there was no active blocking or throttling of Netflix traffic, and in the end Netflix quality has improved by a mechanism that amounts to paid prioritization but does not fall under the FCC’s definition. The FCC has also said there will be “no unbundling,” which is a mechanism that increases competition at the last mile. I make the case here that lack of competition at the last mile is really the core issue that needs to be addressed. The regulations also make no explicit mention of the practice of “Zero Rating” (where bandwidth consumed by certain content providers does not count against the users’ quota). Zero Rating is a much more serious and real anti-

a <https://www.amacad.org/content/publications/pubContent.aspx?d=1263>

Lorem ipsum dolor sit amet consectetur

competitive problem than the ones the FCC has clamped down on (throttling/blocking/paid prioritization). In other words, what the FCC has missed is more important than what it is attempting to fix. Let me walk you through my argument for that claim.

The Changing Internet

Many of us remember the Internet from the days where all packets were indeed “equal” and the commodity that was traded among ISPs and between ISPs and customers was bandwidth. The unit of trade was normalized to bandwidth, and the applications on the Internet were primarily Web and email and all packets could be treated uniformly and charged uniformly as well.

The online world looks remarkably different today and the commerce the Internet generates goes beyond the trading of bandwidth. Trillions of dollars are generated using the Internet, whether we are talking established players such as Google and Netflix or upstart companies such as Uber or Airbnb. While it

is tempting to think of every packet being equal, the revenue each packet generates is different, for example, for a search page (with ads) of Google versus a frame of a Netflix video. This difference needs to be modeled and accounted for.

This change in the economic value of traffic has thrown the economic stability of the Internet out of kilter. It shows up in disputes that are broadly classified as network neutrality issues. Before explaining the specific issues, I provide some background on Cooperative Game Theory.

A Cooperative Game Theory Lens

We looked at the Internet as an economic ecosystem, applying the lens of Cooperative Game Theory—specifically using the tool of Shapley Values. Cooperative Game Theory analyzes coalitions where a group of players (called a coalition) come together and generate some value, V . The solution of a cooperative game is the share of V that goes to individual players. A subset of cooperative games considers so-called convex games, where the

marginal value of the coalition grows with the size of the coalition. Colloquially, this translates to “the whole is greater than the sum of the parts.” The Internet is a good example of a convex game: various folk laws that define the value of the network as a function of the number of nodes as $V(n) = n^2$ (Metcalfe’s law) or $V(n) = n \log n$ (Odlyzko’s law) have shown good agreement with empirical data. Regardless of the precise form of the relationship, everyone agrees on the term “the network effect.” There is general consensus on the “convexity” of the Internet.

Convex games have a set of solutions that are called the “core,” where the share of the value that goes to each player creates an incentive for them to remain as part of the game. If there is a solution outside the core, some players have the incentive to leave the coalition, because they are better off not playing (either they are generating negative value because of costs involved, or they generate more value being on their own or part of some other coalition).

The Shapley value is a specific (and unique) allocation that can be axiomatically derived. One primary axiom that determines the Shapley value is that of “balanced contribution,” which in simple terms means the value a player ends up obtaining is reflective of the contribution the player made. In convex games, this aspect of Shapley values shows up in the following way: the Shapley value solution lies in the center of gravity of the core; that is, it is the most stable of the solutions. It is the solution that keeps all the players as far away from leaving the coalition as possible.

The ISP Settlement Problem

We applied the Shapley value framework to analyze the complex topologies of the Internet and the different players such as ISPs, content providers, and (eyeball) customers. Following the taxonomy introduced by Dave Clark, we classified ISPs into three kinds:

- ▶ Eyeball ISPs that provide broadband connections to customers (such as Comcast and Verizon);
- ▶ Content ISPs that provide connectivity to content providers (we include content delivery networks in our definition), for example, Cogent, Level 3, Akamai; and
- ▶ Transit ISPs that provide global connectivity (such as AT&T, Telefonica, and Tata).

We modeled the revenue collected at the eyeball side as a flat-rate charge per customer (as is the norm for wired broadband), and at the content and transit level as a volume-based charge—which is again the norm. We also modeled the typical customer-provider relationships that exist on the Internet—with eyeball and content ISPs typically buying bandwidth from transit ISPs, and the transit ISPs themselves engaging in settlement-free peering with each other. As a thought experiment, we analyzed whether bilateral trading of bandwidth as a commodity could lead to a value allocation that was in the vicinity of the Shapley value, indicating a stable ecosystem. The idea behind our analysis was not to think of the Shapley value as a prescriptive solution for fair profit sharing, but rather look at the natural mechanisms that existed on the Internet and whether they end up in a stable region or not.

What we discovered through our analysis¹ was:

▶ As long as the flow of traffic on the network remained symmetric and the revenue collected at the two ends (eyeball end and content provider end) remained approximately similar, the bilateral trading in bandwidth resulted in a solution close to the Shapley value, indicating stability. This is reflective of the Internet of 15 years ago. Buying or selling of bandwidth resulted in a stable ecosystem.

▶ When the flow of traffic becomes asymmetric and the revenue generated at one end is imbalanced (on the content side, lump together the revenues that are being generated by Google, Netflix, and Amazon), the bilateral trading of bandwidth as the commodity leads to a solution outside the core. This led us to predict the rise of paid peering on the Internet to push the solution back inside the core. In simple terms, it is a way to transfer some of the value the content providers are generating (for example, Netflix subscription fees) to the ISPs to come up with a net value allocation that is in the core. The implicit trade that occurs now is not one of bandwidth, but that of revenue generated by content. Part of the revenue the customers of an eyeball ISP add to a content provider gets transferred to the ISPs via paid peering.

▶ The location of the solution also depended heavily on the level of competition. The more asymmetry in the level of competition, the less stable was the solution.

Our predictions started coming true in 2013–2014 when Netflix began signing paid peering arrangements with

There is a lot of “competition” on the content side of the Internet, but not enough at the eyeball side.

the major eyeball ISPs in the U.S. They ended up signing these paid peering deals with the top-four major ISPs in the U.S.: Comcast, Time Warner, Verizon, and AT&T. What is also interesting is all these four major ISPs are “local monopolies” in their respective geographical areas. The Netflix dispute with the ISPs has been at the center of the network neutrality debate, and the issue is purely economic. The public peering points between Netflix and the ISPs started experiencing congestion. However, this was not active throttling or blocking (practices the FCC is now explicitly prohibiting). The ISPs could afford to play hardball with Netflix because the customers of the ISP would not discontinue their broadband service if the quality of Netflix streaming was bad. The Internet provides enough other content for it to be worth it for the customer. In other words, there is a lot of “competition” on the content side of the Internet, but not enough at the eyeball side, and that creates an asymmetry predicted and quantified by our analysis.

By signing these paid peering arrangements, the eyeball ISPs are able to monetize the excess bandwidth the content generates and the value allocation solution returns to inside the core. What is also interesting is that while the network neutrality debate has centered on paid prioritization or throttling at the last mile, none of the ISPs were actually doing any of that. The congestion was happening at the peering links of Netflix (or their content delivery networks, Level 3) and the ISPs. By signing a paid peering arrangement, the ISPs have not created a “fast lane” but instead have created a fast “on-ramp” for the Netflix traffic that bypasses the open Internet peering points. The fight was economic, and the resolution was economic as well. There was no suppression of freedom involved—all the entities are interested in making money and that is the primary motivation. This problem arose because the level of competition at the last mile in the U.S. is extremely low, which throws the solution outside the theoretical core. The solution to this should be an approach that increases competition, such as unbundling or municipal broadband, rather than increasing regulation. A lot of places where there is thriving competition at the last mile (for example, Scandinavia and the

Far East) have no real or perceived network neutrality violations, and the Internet service there is fast and inexpensive.

Flat-Rate Versus Volume-Based Charging

A related issue is how ISPs charge for traffic. Customers much prefer paying a flat rate amount, as it reduces their decision making to a once a month “is it worth it?” rather than saddling them with a constant flow of micro decisions for every connection to the Internet. For the ISPs it is bad news because they are unable to monetize excess traffic for flat-rate charging. In the wireless world on the other hand, the norm is volume-based charging (we include monthly quotas plus overage charges in the broad definition of volume based charging). Unsurprisingly, debates about network neutrality do not appear there. It is in the interest of ISPs to have high-volume content such as YouTube or Netflix on their networks, because they can monetize that traffic.

The issue of Zero Rating is a related one that needs a careful examination here. In this scenario, content providers pay for the bandwidth the customers use for their content. This makes the ISPs happy as well, as they get paid for bandwidth—although this practice is truly anti-competitive. If a user gets to access Facebook or Netflix for free, then that user is unlikely to go to a different social network or streaming provider and pay for that bandwidth. For that reason, in the wireless world the big content providers are promoting Zero Rating—Facebook for instance is implementing something called internet.org in certain countries, where bandwidth for content from Facebook and other select sites is not charged to the customer. Zero rating is a much more serious threat to network neutrality than throttling or blocking. And the root cause is again economic. The content providers that seem to be on one side of the fight over network neutrality for the wired world are on the other side of it in the wireless world. Thus, it comes down to how the customer is charged—flat rate or per byte.

Innovation at Risk?

Network neutrality is an issue where the technical details get set aside for an idealistic statement that makes

The network neutrality issue is really about economics rather than freedom or promoting/stifling innovation.

the case that “every packet should be treated equally.” This is simply not true for modern networks—VoLTE is a very good example where voice traffic is prioritized over other traffic on LTE networks, bringing about the first significant improvement in voice quality in many years. According to the FCC ruling, if you are an application that is running with a public IP address, best effort is all you can get from the Internet. If any other kind of prioritized service is requested, it has to be classified as a managed or specialized service. This implies that a VoIP service from a cable provider can get prioritized service, but applications such as Skype and Vonage have to work with best-effort service. Similarly, health applications on smartphones or smart watches are becoming commonplace; under the new rules, an application such as dedicated heart-monitoring would have to come under the classification of specialized services, rather than use the data services already available on their smart devices. This has the potential to slow down innovation on the Internet, because the implication is that any kind of application that requires something more than best-effort service has a barrier to entry placed by regulations.

Differentiated services are routinely implemented in modern networks, and the designers of the Internet protocols foresaw this need by creating specialized Quality of Service bits as well as QoS architectures via a mechanism called DiffServ. The FCC ruling, however, can lead to the potential of the mechanisms remaining unexploited. There could be many other QoS sensitive applications that we cannot even imagine today that

might need to run on the public Internet; we certainly have no idea where the world of Internet of Things would end up. It seems to be a technical risk to disallow any differentiation over the public Internet. The network should be allowed to differentiate but not discriminate, which is an important distinction that needs to be widely understood. Express delivery of packages in the physical world is allowed between all public addresses by all postal carriers, so there is certainly no reason to limit that capability on the Internet.

Conclusion

Our analysis over the past several years has revealed the network neutrality issue is really about economics rather than freedom or promoting/stifling innovation. This issue arose because the Internet has slowly transitioned from an entity that was primarily used for research purposes to the centerpiece of the modern economy. However, the reaction to the economic disputes that have naturally arisen has been to introduce regulations risking that they could potentially slow down if not completely limit innovation on the Internet. On the flip side, the lack of focus on the issue of competition can lead to a further strengthening of incumbent monopolies, which reduces the incentive for innovation. Making a policy decision without fully understanding and accounting for the science behind it is risky, from both an economics perspective and a networking one. The way forward lies in increasing competition at the last mile, rather than regulations that might impose technical limitations on the (future) functioning of networks. Network neutrality is the symptom; the real issue is lack of competition at the last mile. ■

Reference

1. Ma, R.T.B., Chiu, D., Lui, J.C.S., Misra, V., and Rubenstein, D. On cooperative settlement between content, transit, and eyeball Internet service providers. In *Proceedings of the 2008 ACM Conference on Emerging Network Experiment and Technology (CoNEXT 2008)*, Madrid, Spain, December, 2008.

Vishal Misra (misra@cs.columbia.edu) is a professor of computer science at Columbia University.

Copyright held by author.