intermedia
www.iicom.org
May 2013 Volume 41 Issue 2

By Christopher Yoo

The non-monolithic Internet

Introduction

In recent years, policy debates have exhibited a tendency to talk about the Internet as if it were a single, unified phenomenon to which everyone must have non-discriminatory access. I would like to examine this proposition critically and explore the policy implications that might follow if it turned out not to be true. Exploring the ways that the Internet has already deviated from this "single Internet" vision in the past and is likely to continue to do so in the future may provide insight into how best to shape Internet policy.

IP-based networks that bypass the public backbone

Despite the fact that discussions often somewhat casually treat the Internet as a single entity, the Internet is widely recognized by those familiar with its operation to be a "network of networks", composed of many smaller networks interconnected together. All of those networks exchange information through a uniform standard known as the Internet Protocol (IP), which represents the glue that binds the components into a larger coherent whole. In addition, Internet traffic has traditionally traversed backbone providers that exchanged information at public exchange points. When this was the case, networks exchanged traffic in a largely uniform manner.[1]

Over time, an increasing amount of traffic has begun to deviate from the traditional pattern. Most importantly for our purposes, some firms rely on the Internet Protocol, but transmit their data over proprietary networks. For example, during the Internet's early days, the acceptable use policy imposed on the NSFNET (the National Science Foundation Network) prohibited uses unrelated to research and education. This prompted the creation of private backbones such as PSI, UUNET, and CerfNET, which operated in parallel with the NSFNET backbone without interconnecting with it.

In later years, firms began providing IP-based voice services without traversing the public Internet in order to ensure improved security or greater guaranteed quality of service, which caused these services to be called "IP telephony" or "voice over Internet protocol" (VoIP) instead of Internet telephony. In addition, a large proportion of IP-based video passes via managed services to distributed caches without touching the public backbone (Yoo 2010a). A prominent example is Netflix's new content delivery network, Open Connect. In short, a significant amount of IP-based traffic is travelling over various proprietary networks that bypass the public backbone altogether.

Hybrid networks

A large amount of additional IP-based traffic employs hybrid systems that rely on proprietary or non-IP based technologies to form part of the connection. For example, third-generation wireless networks employed hybrid systems that used legacy, non-IP based technologies to connect from the Internet gateway to the end-user device.[2]

Still other networks are moving away from exchanging data on the basis of IP in favour of technologies that support greater functionality. One of the most popular of these is Multi-Protocol Label Switching (MPLS). Instead of routing based on IP addresses, MPLS adds a label to the front of each packet and routes on the basis of that label. In addition, each flow (known as a forwarding equivalence class) is assigned a specific path through the network. Information about the label and the associated route are propagated to other MPLS-enabled routers.[3] Because labels are shorter than IP addresses, routers can direct traffic more rapidly. The fact that the route for a particular flow is defined in advance gives end users greater control over security. In addition, MPLS can support load balancing simply by dividing traffic between the same two endpoints into two separate forwarding equivalence classes and assigning them different paths.

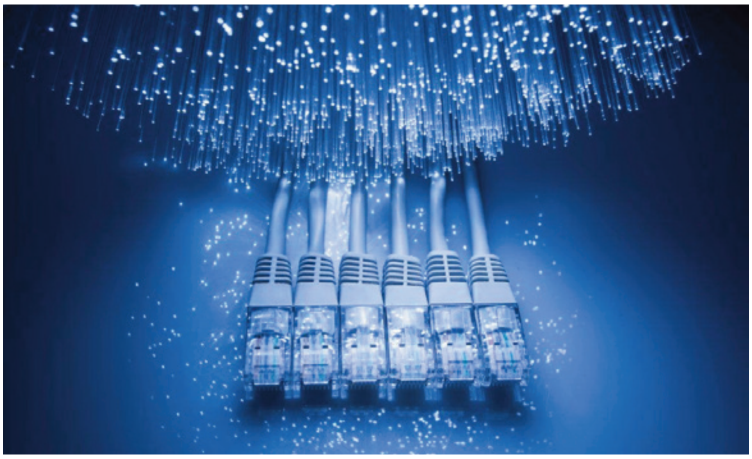

Most importantly, in determining the particular path that a flow will travel, the MPLS router can match the quality of service demanded with the resources available along possible paths.[4] MPLS can also recover more quickly from route failure. While once used exclusively within a single network, MPLS is now being employed across multiple networks, although doing so requires elaborate coordination between the networks to exchange the necessary information about the labels.

In short, a significant amount of data traffic relies on proprietary or hybrid networks rather than the public Internet. A leading industry study estimates that more than 22% of all IP-based traffic did not rely exclusively on the public Internet in 2010, and it forecasts those numbers to remain more or less stable for a few years before tailing off to 16% by 2016.[5] Most of these networks require the negotiation of private interconnection agreements and engage in extensive network management and prioritization.

Positing the policy challenge as a mutually exclusive choice between managed and unmanaged networks may represent a false dichotomy. The routers within the core of the network are becoming increasingly programmable so that they can be dynamically reconfigured from the default setting of providing traditional Internet service to provide a dedicated circuit instead.

One leading example is Internet2's Interoperable On-demand Network (ION), which allows researchers to set up temporary or long-term dedicated point-to-point optical circuits to support large data transfers and other bandwidth-intensive applications.[6] This represents a sea change from Internet2's previous position disfavouring network management as a solution.[7]

The ability to reconfigure routers dynamically illustrates the problem with framing the decision as a choice between two diametrically opposed alternatives. Instead, the competing visions of a single, openly accessible network on the one hand and a completely Balkanized universe of non-interconnected private networks on the other represent polar extremes on a spectrum of possible policy responses. Only by understanding the countervailing considerations can one understand the forces that determine where along that spectrum social welfare would be optimized and what types of changes in the economic and technological environment might cause the optimal balance point to change.

Private networking as an exit option

The assumption that the Internet must remain open and universally interconnected also overlooks the fact that the imposition of government regulation rarely results in a stable equilibrium. Instead, the ultimate impact of such regulations can only be understood after the series of reactions and counter-reactions that the regulations are likely to engender are taken into account.

Consider the likely impact of mandating interconnection and nondiscrimination on a best-efforts basis. Providers who need quality of service guarantees would have strong incentives to bypass the public Internet - and the accompanying regulations - simply by shifting their resources to private networks. The problem is that shifting to dedicated networking resources eliminates one of the central efficiencies created by the Internet by preventing multiple users sharing resources. In particular, network operators who need quality of service assurances will generally find it most efficient to offer their bandwidth to others on a secondary basis, available only when the primary owner does not need it.

Those who advocate mandating open access in the name of preserving a unified Internet thus must consider the possibility that such a mandate might create incentives toward greater fragmentation.

In addition, restricting network owners' ability to share surplus capacity threatens to increase the cost of broadband where it is available, while at the same time worsening the digital divide by reducing the geographic areas in which such services can break even.

Conclusion

Any assessment of claims that unity and openness are critical features to the Internet's success must grapple with the extent to which the Internet has been disunified in the past and in the present. In fact, the Internet is not and never has been a monolith. It is a collection of standards with respect to which people have always had the option to opt in or opt out, either partially or completely. Indeed, many have and will continue to do so.

Moreover, the unqualified manner in which this "single Internet" claim is often advanced obscures the fact that important policy decisions typically involve a trade-off between competing considerations. Mandating open interconnection implicitly presumes that the exclusive source of value to end users is raw increases in network size. Framing the issue in this manner fails to consider that end users typically place a premium on being able to reach a small number of locations and run a discrete number of applications.[8]

The policy analysis must also take into account that end users are likely to maintain more than one connection and that any attempt to mandate open access is likely to provoke a series of reactions and counter-reactions that may frustrate the goals of the initial regulation.

On a broader level, policymakers would benefit from taking seriously the possibility that the days of a "one size fits all" approach to Internet regulation may well be over and that looking backwards for the lessons of the past may not always be the best way to promote future success.[9] As anyone can attest who experienced how quickly AOL keywords shifted from the critical way to access customers to near-complete irrelevance, the technological and economic environment surrounding the Internet is constantly undergoing rapid, dynamic change. The policies developed for a world dominated by PCs using cable modem or DSL services in which the browser was the critical platform may no longer be the right framework for a world increasingly dominated by smart phones attached to wireless broadband networks in which the critical platform is now the app store and the wireless operating system. The growing heterogeneity of the technologies, end user demands, and business relationships underlying what is now often referred to as the Internet ecosystem may require reframing the issues in a fundamentally different manner.

Christopher S Yoo is John H Chestnut Professor of Law, Communication, and Computer & Information Science and Founding Director of the Center for Technology, Innovation and Competition, University of Pennsylvania. This contribution is adapted from Christopher S Yoo (2012) "Internet policy over the next five years: Does one size still fit all?" pp 51-65 in Randolph J May (ed.) *Law and communications policy in the digital age: The next five years.* Durham: Carolina Academic Press.

---

[1] Christopher S Yoo (2010) "Innovations in the Internet's architecture that challenge the status quo" *Journal on Telecommunications and High Technology Law* **8** (1) 79-99.

[2] Christopher S Yoo (2010) "The changing patterns of Internet usage" *Federal Communications Law Journal* **63** (1) 67-89.

[3] Eric C Rosen, Arun Viswanathan & Ross Callon (2001) "Multiprotocol label switching architecture" IETF Request for Comments 3031.

[4] William Stallings (2001) "MPLS" *Internet Protocol Journal* **4** (3) 2-14.

[5] Cisco (2012) *Cisco Visual Networking Index: Forecast and Methodology, 2011-2016.* San Jose: Cisco Systems, Inc.

[6] Internet2 (n.d.) "Internet2 ION: user-provisioned dedicated circuits reserved in advance or on demand" available at http://www.internet2.edu/ion/

[7] US Senate (2006). *Net neutrality: Hearing before the Senate Committee on Commerce, Science, and Transportation,* 109th Cong., 2d Sess. 63-68.

[8] Christopher S Yoo (2012) "When antitrust met Facebook" *George Mason Law Review* **19** (5) 1147-62.

[9] Christopher S Yoo (2012) *The dynamic Internet: How technology, users, and businesses are transforming the network.* Washington, DC: AEI Press.