The Contingent Internet

David D. Clark

Abstract: The Internet is so omnipresent and pervasive that its form may seem an inevitability. It is hard to imagine a "different" Internet, but the character of the Internet as we experience it today is, in fact, contingent on key decisions made in the past by its designers, those who have invested in it, and those who have regulated it. With different choices, we might have a very different Internet today. This paper uses past choices made during the emergence of the early Internet as a lens to look toward its future, which is equally contingent on decisions being made today: by industry, by governments, by users, and by the research community. This paper identifies some of those key choices, and discusses alternative futures for the Internet, including how open, how diverse, how funded, and how protective of the rights of its users it may be.

DAVID D. CLARK, a Fellow of the American Academy since 2002, is Senior Research Scientist at the MIT Computer Science and Artificial Intelligence Laboratory. He has been involved in the design of the Internet since the mid-1970s and is a member of the Internet Hall of Fame. His recent policy publications include a chapter in *Trust, Computing, and Society* (ed. Richard H. R. Harper, 2014), and articles in the journals *Telecommunications Policy* and *Journal of Information Policy*.

Is it possible that the Internet might never have happened? Is it possible that, in a parallel universe where the Internet's inventors had pursued different careers, we could be without a network that links all of our computers together? That we might have "personal computers" that were truly personal, not connected to the larger world unless their contents were copied to disk and mailed?

Actually, that alternative outcome is highly improbable. The Internet was in some respects a creation of its time: in the 1960s, the idea of a global network for computers was "in the air." A visionary of the time, J. C. R. Licklider, had already predicted teleconferencing, information sharing, instant messaging, online tax preparation, offshoring, and the potential for a digital divide.[1] However, at the time of the Internet's launch, there were competing conceptions for how to build a "computer network." Our alternate universe is not without the Internet, but rather is with a very *different* Internet.

This possibility may itself seem surprising: the Internet today is so omnipresent, so much a fixture of our lives that it seems almost as if it "had to be that way." What might an alternate Internet have looked like? This is an important question, because to recognize that there were multiple options for the early Internet, and that the Internet as we know it is contingent

on decisions that could have led to different outcomes, is to recognize that the future of the Internet is itself contingent. Society will meet forks in the road that will determine the future of the Internet, and recognizing these points and discussing the alternatives, rather than later looking back and wondering if we chose the right path, is an opportunity we cannot forego.

The Internet is a "general purpose" network, designed for a variety of uses. It is suited to email, watching video, playing a computer game, looking at Web pages, and myriad other applications. To an Internet engineer, the Internet is the system that moves data, and the applications (like a Web browser, which users might lump into the larger concept of "Internet") run on top of that data-transport service. This modularity, and this generality, seem a natural way to structure a network that hooks computers together: computers are general-purpose devices; since the Internet hooks computers together, it too ought to be general. But this idea was quite alien to the communications engineers of the early-Internet era, who largely worked for telephone companies. They asked what was to them an obvious question: how can you design something if you don't know what it is for? The telephone system was designed for a known purpose: to carry telephone calls. The requirements implied by that purpose drove every design decision of the telephone system; thus, the engineers from the world of telephone systems were confounded by the task of designing a system without knowing what its requirements were. The early history of the Internet was therefore written by people who came from a computing background, not a classical network (telephone) background. Most computers are built without a singular purpose, and this mind-set drove the Internet's design.

But this generality has a price. The service the Internet delivers is almost certainly not optimal for any particular application. Design for optimal performance and design for generality are two distinct objectives. And it may take more effort to design each application in a general network than in a network that was tailored to each application. Over the decades of the Internet's evolution, there has been a succession of dominant applications. In the early years of the Internet, the Internet was equated to email, and to ask people if they were "on the Internet" was to ask if they had an email address. Email is an undemanding application to support, and if the Internet had drifted too far toward exclusively supporting it (as was happening to some degree), the Web might not have been able to emerge. But the Web succeeded, and its presence as a complement to email reminded engineers of the value of generality. But this cycle repeats, and the emergence of streaming audio and video in the early 2000s tested the generality of an Internet that had drifted toward a presumption that now the Web, and not email, was *the* application. Today, streaming, high-quality video drives the constant reengineering of the Internet, and it is tempting once again to assume that we know now what the Internet is best suited for, and optimize it accordingly. The past teaches us that we should always be alert to protect the generality of the Internet, and allow for the future even when faced with the needs of the present.

There is another aspect of generality: the applications that run over the basic transport service of the Internet are not designed or distributed by the same entity that provides the basic data-transport service. This characteristic has been called the "open" Internet, and again, this separation made sense to a computer engineer but did not fit conceptually with the telecommunication engineer. The telephone company installed that wire to your house to sell you telephone service, not to enable some other company to sell you theirs. From the telephone com-

pany's perspective, it is expensive to install all those wires, and how could they get a reasonable return on investment if they were not the exclusive service provider?

In the early days of the Internet, the only way to access the Internet from home was to use a modem to make a dial-up connection to an Internet service provider (ISP). A residential user paid the telephone company for the telephone service, and then paid the ISP for providing access. This seemed then like a minor shift in the business model of the telephone companies. But as the possibility of expanding broadband services to the home emerged in the 1990s, the corporate resistance to an open platform became quite clear. One telephone executive explained to me at the time: "If we don't come to your party, you don't have a party. And we don't like your party very much. The only way you will get broadband to the home is if the FCC forces us to provide it."

That was a fork in the road, and the Internet certainly might have taken another path. In fact, the force that led the Internet toward residential broadband was, to a considerable extent, the emergence of the cable television industry as a credible and competitive provider of high-speed residential Internet.

We continue to see echoes of this tension between the Internet as an open platform for third-party applications and broadband access as an expensive investment that should work to the advantage of its owner. The current debates around the concept of "network neutrality" are at their heart about whether broadband providers should be regulated to provide a neutral, open platform for third-party services, or if they have the right to define the services they offer (and perhaps favor) over the infrastructure they invested in building.

Another consequence of generality is that the data-transport layer of the Internet has no concept of what the application is trying to do (as opposed to the design of the tele-

phone system, which at all levels reflects the centrality of the telephone call). If the design of the Internet required that the network understand what the application were doing, deploying a new application would require its designer to somehow modify the core of the network to include this knowledge. To the early designers, this was a fork in the road down which they did *not* want to go. If an application designer had to alter the network before deploying a new application, this would both complicate the process of innovation and create potential for the network to block one or another application.

The Internet has been called the *stupid network*, the telephone system being the *intelligent network*; the open-design approach of the Internet makes perfect sense – that is, until things go wrong. If the network itself is impairing the operation of an application, the network cannot always detect or correct this. The network may be able to detect that one of its components has failed, but more complex failures may go undetected, leaving frustrated users who can see that their application is not working, but who have no remedy available to them. Had we taken the fork in the road that enabled the network to know more about what each application was trying to do, the network might have been less supportive of easy innovation, but might also have been less frustrating to use when unexpected problems inevitably arose.

Finally, the division of responsibility between the provider of the data-transport service and the provider of the application means that responsibility for core requirements like security is divided among several actors. This both makes the objective harder to achieve and adds incentive to delegate the task to another party. In this way, the design decisions that shaped the Internet as we know it likely did not optimize secure and trustworthy operation.

David D. Clark

These design choices led to differences in the technical character of the Internet, but many choices also led to particular outcomes in the industrial structure of the Internet ecosystem. When we made design decisions about system modularity in the early Internet, it was not entirely clear to us that our design was both a technical structure and an industrial structure. Some of the early network pioneers, though, certainly did understand this. In the 1970s, there was a substantial debate between advocates of two sorts of networks: *datagram* and *virtual circuit*. Datagram networks have a simpler core, with more functions shifted to hosts at the edge. Virtual-circuit networks have more function in the core of the net, and thus more power and control shifted to the network operator. The Internet is a datagram network; the ARPANET (Advanced Research Projects Agency Network) that preceded it was a virtual-circuit network.

One of the most vocal advocates of the datagram approach was the French computer systems designer Louis Pouzin, who was building a datagram network called Cyclades at the same time that the Internet was taking shape. In 1976, he published a paper that reached the following conclusion:

> The controversy DG vs. VC in public packet networks should be placed in its proper context.
>
> First, it is a technical issue, where each side has arguments. It is hard to tell objectively what a balanced opinion should be, since there is no unbiased expert. This paper argues in favor of DGs, but the author does not pretend being unbiased. Even if no compromise could be found, the implications would be limited to some additional cost in hardware and software at the network interface. So much resources are already wasted in computing and communications that the end result may not be affected dramatically.
>
> Second, the political significance of the controversy is much more fundamental, as it

signals ambushes in a power struggle between carriers and computer industry. Everyone knows that in the end, it means IBM vs. Telecommunications, through mercenaries. It may be tempting for some governments to let their carrier monopolize the data processing market, as a way to control IBM. What may happen, is that they fail in checking IBM but succeed in destroying smaller industries. Another possible outcome is underdevelopment, as for the telephone. It looks as if we may need some sort of peacemaker to draw up boundary lines before we all get in trouble.[2]

Pouzin saw the battle over control of the global network as a battle between the computer industry and the telecommunications industry. At the time, the computer industry was dominated by huge players like IBM, giving shape to Pouzin's "Battle of the Titans." IBM was a vertically integrated corporation, just as the telephone companies were: if a firm got its hardware from IBM, it likely got its software from IBM as well. Pouzin may not have foreseen the coming shift in the computer industry to support more open-hardware platforms, but he clearly saw different technical decisions as shifting the balance of power from one industry sector to another.

In contrast to the Internet, Pouzin's Cyclades network was ultimately unsuccessful. Its failure is (speculatively) attributed to the hostility and resistance of the French PTT (postal, telegraph, and telephone government unit).

One of the lessons of the past is that the users of the Internet are an active force in defining what the network is, both by choice of application and by the creation of unexpected applications never anticipated by network engineers. This trend continues today with the success of user-created programs for peer-to-peer music sharing, for example. Sometimes users take the network down a fork in the road that the designers

did not contemplate, or perhaps had even dismissed. In the very early days of computer networking, the designers were focusing on remote access to the expensive, high-power computers of the time. One of the early network engineers asserted that message services were "not an important motivation for a network of scientific computers."[3] Of course, users proved him wrong, flocking to email en masse.

The second, related lesson is that the open character of the Internet is what allows this sort of user-driven evolution to take place. In a more structured and vertically integrated vision of computer networking, the network provider might not even choose to offer an email application. The Internet, by its structure, is amazingly open to exploration by users and third-party innovators. This benefit, though, was clearly contingent on earlier design choices.

Is the core technology and structure of the Internet set for the indefinite future? Or are there further forks in the road that might change the basic character of the Internet? One important and ongoing debate concerns the extent to which encryption should be used by default to protect communication between users from observation (and modification). If the goal is privacy of communication among communities of users, encryption is a powerful tool. But encryption thwarts the goals of many other actors: the intelligence community benefits greatly from being able to spy on content, and in some nations, this right of spying is not subject to debate. If the Internet were to move to a posture of "encryption by default," would certain nations opt out of the public Internet as we know it and essentially create a separate, state-controlled network for their citizens? Further, Internet service providers find encryption between communicants problematic, since it prevents them also from seeing what their users are doing. If the Internet were still the totally open,

neutral platform that only moved data from sender to receiver, such peeking would seem unnecessary; but that interpretation of today's Internet is an oversimplification. Operators claim that they need to see what users are doing in order to optimize their experience – which may at times be true – and to selectively influence what they are doing. Many of these interventions by Internet service providers – such as the modification of data in transit to insert advertisements – have been met with protest, but other interventions – like reformatting data to fit content onto the small screen of a mobile device – are more easily justified by users. The tussle over the use of encryption is only the latest chapter in the struggle between the providers of the data-transport service of the Internet, the providers of higher-level services and applications, the users, and the state system for control of the user experience.

The experience of using the Internet is becoming more diverse. While the data-transport service is more or less uniform across the globe, that service does not define the user experience. The experience of the user is defined by the applications that run "on top of" that service. If different applications are available or are preferred by users in a given region, the resulting Internet experience will, in turn, be different. Perhaps the most obvious example of this today is the Internet experience in China, where the state government has blocked access to many of the applications that define the Western user experience, such as Facebook and Twitter. But there are domestic equivalents to these services within China that make the Chinese Internet experience a vibrant space of interaction, even if it is heavily policed. Still, there is no easy way for a Chinese and American to "friend" each other using Facebook.

The global community will need to decide the extent to which we fight against this diversification of the user experience. Some of the early Internet visionaries con-

David D. Clark

ceived the Internet as a platform for global discourse and a vector for a global civil society. Diversification of the Internet experience would seem to erode that vision. On the other hand, differences of language, culture, and norms are real, and would suggest that as the Internet matures, the experience in any region would evolve to conform to those factors. The generality of the platform does not mean that everyone has to use it the same way. In fact, the generality, combined with the ability of users to vote (via their usage) for the applications of their choice, almost makes regional diversity inevitable. If the nations of the world were to push for some sort of global alignment of regulation and incentive, it would likely lead to a more homogenous but less satisfactory Internet experience. Perhaps the ideal is an Internet that accepts the diversity of experience for most users, but permits interaction on a global level among users who seek it. We should urge China not to block Facebook (which they view as a threat to regime stability) but we should at the same time accept the outcome that most Chinese users prefer their domestic alternatives.

Another critical issue that will shape the future Internet is the poor state of Internet security. We hear almost daily about theft of data, computers corrupted with malicious software (malware), cyber-crime, and many other breaches of the security of both users and service providers. This state of affairs could play out in a number of ways. One is that the current state of insecurity persists, which might eventually prevent sensitive or important transactions from taking place online. In this way, poor security could be a barrier both to the uptake and utility of the Internet. An alternate future is that in the attempt to improve Internet security, the Internet mutates to hold users more accountable for their actions. For a number of reasons, the original design of the Internet did not include any mechanisms to deal with identity management. It was understood,

if imperfectly, that different applications might call for differing degrees of accountability; for example, while a transaction between a customer and his or her bank calls for strong mutual verification of identity, another user may not be comfortable searching for information on AIDS, for example, if the query is not anonymous. Defining which actors can issue globally trustworthy identity credentials is, of course, another challenge to uniform identity management. These concerns remain valid; but at the same time, the pressure to improve our ability to hold users accountable and deter malicious behavior may push toward stronger identity tracking. In many nations today, one must provide a national identity number to use the Internet; respect for anonymous action may only be a local preference, which could erode under global pressure for accountability.

The explosion of mobile devices signals another inflection point for the future of the Internet. The traditional industry narrative pitted application designers against the Internet service providers: application designers wanted an open, neutral platform on which to innovate and the Internet service providers wanted control over the services offered in order to effectively monetize the user experience. Computer manufacturers were seen as neutral in this dynamic, lacking a business model that bundled proprietary applications on the device. But the trajectory of the mobile device is very different: the makers of smart phones show a much greater interest in shaping (and monetizing) the user experience. Apple charges a fee (currently 30 percent) for the sale of a paid application for the iPhone or iPad through their app store. The regulators in the United States would be quick to intervene if ISPs tried to charge a fee to customers for using a specific app, but so far there has been little criticism of device makers doing exactly that.

The interplay of device maker, ISP, and application creator and provider is partic-

ularly interesting in developing nations. Getting developing states online is socially desirable; the power of connectivity to improve the conditions of citizens is evident. But what strategies are acceptable in pursuit of this goal? Here, Facebook has launched a clever scheme: they developed a stripped-down version of the Facebook application (called 0.facebook.com; the zero implies zero cost) and have negotiated an arrangement with mobile service providers in many developing countries ensuring that use of the application is free to users, not counting against any data quotas. By making use free – and in some cases even arranging for a discount on the device – the uptake of Internet in the developing world may increase. But as a consequence, a generation of users will equate the Internet not with sending email, not with searching the Web, but with using Facebook. This is not a hypothetical outcome; surveys suggest it is already happening.[4] Is this degree of corporate capture acceptable? And what growth potential does it limit? Of course, this is an issue that each country will decide for itself through domestic regulation (or deregulation, for that matter). But again, such decisions will likely diverge the character of the Internet experience as it evolves across the globe.

The design alternatives I have described may seem to concern principally the user experience, but at a deeper level, they are struggles over control. The previous examples illustrate tussles between ISPs, application designers, device makers, and governments (among others) over control of the Internet. Depending on how the balance of power among these actors evolves, we may see different outcomes with respect to deployment, openness, innovation, and user experience. The immense power of private-sector actors is notable; society has largely left the future of the Internet in the hands of far-reaching, profit-seeking entities. In the United States, there is a tendency to put

our hope in competition, as if with enough competition the market will converge on what users prefer. Sadly, however much we wish, there is simply not much competition to build residential broadband access networks; the investment and risks required to become competitive are too great. And at the layer of applications, there seems to be a recurring tension between building an application that attracts users and developing an application that makes money, perhaps by capturing information about the user to be used in more selective (and thus more pricey) advertisements. Competition cannot discipline this behavior if the behavior alone allows commercial providers of applications to make money.

Thus, another fork in the road is how the Internet will be paid for. Today, aside from some public-sector money that supports specific challenges like rural deployment, there are only two important sources of money to pay for the Internet: the fees we pay as service subscribers, and advertising. Internet users today pay for broadband access, and they pay for an assortment of applications and services, including streaming video, online games, and music services. But advertising pays for the "free" Internet experience, those websites that cost users nothing to visit. And advertising dollars will not grow without bound. Spending on advertising can only be a fraction of commerce, or ecommerce, as is most common on the Internet. In 2013, $42.8 billion was spent on Internet advertising in the United States.[5] The total monthly advertising expenditure per household with broadband access (eighty-eight million households in 2013) is about $40.[6] In other words, all of the advertising-supported Internet content is fighting over an amount that is smaller than the average monthly cost of broadband access. Could the Internet experience stall because we run out of advertising dollars?[7]

Online advertising will indeed grow as it cannibalizes traditional TV advertising. But

David D. Clark

what if a new "Internet experience economy" arises, in which users pay a small amount for access to a broad spectrum of Internet applications that do not track their usage and do not run advertisements? Users are largely accustomed to paying for both premium apps on mobile devices and music and video. Perhaps there will be a shift in how they pay for access to Web content. And were this to happen, who would control that payment ecosystem?

The next major tussle over control involves the governments of the world. The private sector has a common set of motivations: be profitable, grow, survive. Governments have a range of concerns: national security (which can include regime stability), law enforcement, taxation, control of "unacceptable" content, and protection of the rights of powerful private-sector actors (such as protection of copyright), among many others. Different countries have different priorities, different laws, and different approaches to governance. In some cases, these priorities put them at odds with the private sector that, in most countries, has a dominant influence over the character of the Internet. In late 2012, at the World Conference on International Telecommunications in Dubai, the International Telecommunications Union (a division of the United Nations) proposed an international treaty that would give it the right to regulate international interconnection in the Internet. This idea was supported by a number of powerful nations, but it failed to gain traction. However, this preference for state control over important aspects of the Internet will probably continue to grow in certain quarters.

But in all of this contention over the future of the Internet, there is one set of actors that has faded from view: the federally funded research community that designed and built the Internet. From one point of view, this trajectory is proper: they did their job, the commercial world has taken over, and the Internet is now an engine of economic innovation. But from another point of view, there might be a richer, more diverse set of uses for the Internet if nonprofit actors were motivated and supported to develop "noncommercial" applications. As the proliferation of apps for mobile devices suggests, it is not hard to launch a new application today. Perhaps one way to pick among future alternatives for the Internet is for interested parties to vote with their dollars, funding the development of applications that are not motivated by the pursuit of profits, but by interest in civic, cultural, or political participation.

It is possible that if we leave the future of the Internet in the hands of powerful private-sector players, we will get the outcome we want. It is possible that if we allow the governments of the world to make decisions that shape the future of the Internet, we will get the future we want. But there is great risk in being passive about the Internet's future; it may simply be too important to leave either to the forces of commerce or the mechanisms of global politics. Perhaps the most important question is how the voice of the users of the world can be injected into the decisions that will shape the future of the Internet.

ENDNOTES

1 J. C. R. Licklider and Robert Taylor, "The Computer as a Communication Device," *Science and Technology* (April 1968).

2 Louis Pouzin, "Virtual Circuits vs. Datagrams: Technical and Political Problems," *Proceedings of the June 7 – 10, 1976, National Computer Conference and Exposition* (New York: Association for Computing Machinery, 1976), 483 – 494.

3 Lawrence G. Roberts, "Multiple Computer Networks and Intercomputer Communication," *SOSP '67 Proceedings of the First ACM Symposium on Operating Systems Principles* (New York: Association for Computing Machinery, 1967).

4 Leo Mirani, "Millions of Facebook Users Have No Idea They're Using the Internet," *Quartz*, February 9, 2015, http://qz.com/333313/milliions-of-facebook-users-have-no-idea-theyre-using-the-internet/.

5 Interactive Advertising Bureau and PricewaterhouseCoopers, *IAB Internet Advertising Revenue Report: 2013 Full Year Results* (New York: Interactive Advertising Bureau and PricewaterhouseCoopers, 2014), http://www.iab.net/media/file/IAB_Internet_Advertising_Revenue_Report_FY_2013.pdf.

6 National Telecommunications and Information Administration, United States Department of Commerce, "Household Broadband Adoption Climbs to 74.2 Percent," June 6, 2013, http://www.ntia.doc.gov/blog/2013/household-broadband-adoption-climbs-724-percent.

7 Moreover, how will the growth of programs like Adblock, which allows its two hundred million users to surf the Web without seeing most conventional ads, affect advertising revenue?