

What is the Internet?

David Weinberger, Ph.D

self@evident.com

Harvard Berkman Center for Internet & Society

March 1, 2010

Creative Commons Attribution-Noncommercial 3.0

Even within a single culture, "What is the Internet?" is a surprisingly difficult question. I will look at what it means to answer that question, initially from within my American standpoint, and then will ask if aspects of the Internet cross cultural boundaries.

What the Internet is not and what it is

The most common and most important way of going wrong about the Internet confuses the Internet with the technology that implements it. The Internet according to this view consists of the routers, the wires, the satellites, and the servers that deliver the internet to us.

This is such an easy mistake to make that we ought to count it as a temptation: It takes an effort of will not to think of the Internet this way. After all, we use the Internet by having equipment in our hands that connects to other equipment. We pay for that equipment. The pieces we hold in our hands are ours. Taking the Internet to be the tangible equipment of the Internet is an easy mistake to make.

It is also the mistake that the commercial interests that deliver Internet access in America encourage us to make. They have invested heavily in laying cables and installing routers. They have a strong financial interest in viewing the Internet as something that they've built and that they've delivered. They believe they own the Internet, or at least their portion of it.

It seems they're on good grounds, for they have spent a great deal of money buying and deploying technology. They do indeed own the wires they've laid and the rooms full of computing equipment they've installed. They can choose to sell it, or even to destroy it,

all without having to get anyone's permission. They have full property-owner rights over that equipment. Nevertheless, that equipment is not itself the Internet.

It's easy to show that the Internet is not in fact the equipment that instantiates it. First, we could replace all that equipment and the Internet would not be changed. More important, we could come up with entirely new types of equipment and run the Internet over it -- including using carrier pigeons.¹ The great strength of the Internet is that it is independent of all physical implementations, and thus can be implemented using many different physical infrastructures.

This makes the Internet different than the telephone network. The telephone network consists of the equipment that implements it. The Internet is something different: It is an *agreement*.

The idea behind the Internet in the first place was to make it easy to integrate computer networks. The Internet is an inter-network of networks. To add a new network to the Inter-network, only a few rules have to be agreed to. The Internet consists of those agreements, or, to put it technically, a set of protocols.

These protocols do not specify what sort of content can go over the network, or who owns the content, or how who gets to participate. The protocol simply says that if you want to connect your network to the inter-network, you have to agree to move packets of information toward where they're headed without interfering with them.

Now, of course it's far more complicated than that. There are agreements about how to move the packets, how long they have to bounce around before you can let them die, and the like. But, the fundamental agreement -- the one that has made the Internet the global network of networks -- was brilliant in its simplicity: You can join if you agree to move packets closer to where they're headed.

The Internet is an agreement. Any communications network that agrees can join. That is how the Internet grew so quickly. It is also what makes it so resilient. Even more impor-

¹ http://www.techdigest.tv/2007/12/le_web_3_live_t.html#more

tant, the fact that at the technical layer the Internet is a simple agreement has had the most profound effect on the Internet at the social and cultural layers. (Whether technology determines its social and cultural effects is a question we will face later in this paper.)

The brilliance of stupidity

The simplicity of the Internet was a collective act of genius. Previous communication networks had assumed that the best network is the one that provides the best services. For example, the best telephone network is one that provides integrated answering machines, automatic call backs when phones are busy, and caller ID so users can see the phone number of the person who is calling them. Each of these services makes the network a little more valuable by making it a little smarter.

But, each of these services also makes the Internet more appropriate for some particular task. The features the telephone system providers have added over the time have made the telephone system a better network for making telephone calls. Likewise, when the cable TV system adds new features, it becomes better for watching television. But, as one of the creators of the Internet, David P. Reed, has said², every optimization of a networked de-optimizes that network for other uses: A network optimized for carrying telegraphic information is de-optimized for carrying video signals.

Reed was one of the authors of a seminal presentation, “End-to-End Arguments in System Design” in 1981.³ It makes the case for the end-to-end design that was adopted as a basic principle of the Internet’s architecture. An end-to-end architecture connects nodes (or “ends”) as directly as possible, with as few built-in services as possible. This is counter-intuitive, for those services have value, so surely a network with those services would have more value than one without them. But, the paper makes the case that over

² During an interview I conducted in 2002.

³ <http://www.reed.com/Papers/EndtoEnd.html> “End-to-End Arguments in System Design.” J.H. Saltzer, D.P. Reed and D.D. Clark. Also see David Isenberg's “Rise of the Stupid Network” <http://www.isen.com/stupid.html>

the long term, those services are best provided by the ends themselves. For example, if you were designing the Internet, you might reasonably assume that since many people will want to use a search engine, search ought to be built into the network as a service from the beginning. Likewise, you might reasonably assume that people will want some security built in so that they can do trusted transactions. You would be right about the value of those services, but if you had built them into the network itself, we now would be using the best search and security services built in 1981. Far better to allow anyone with a good idea to offer services in an open, competitive market.

The standard example is a good one: It took the telephone network decades to come up with faxing as a service. The creation of this service required telephone company engineers to think it through, telephone executives to decide on the business model, and the marketing department to let people know about it. It was a highly successful offering, but the telephone companies' pace in introducing such services is glacial. On the Internet, anyone with a new idea for a service can put it up on a site -- an "end" -- and count on the Internet providing the same transportation of packets as it provides to gigantic corporations. The end-to-end principle is responsible for the Internet becoming the greatest engine of innovation in human history. The authors of the end-to-end principle are arguably responsible for the greatest creation of wealth worldwide since the invention of money.

This is for a simple reason that sets the Internet apart from other technology. Telephones are for talking with people. Cars are for driving. Chairs are for sitting. What is the Internet for? It is for anything that can be done by moving packets of information around. These packets can carry information that computers on the end points will interpret as text, as graphics, as voice, as movies, as medical information, as signals for turning on the dishwasher before you arrive home, as ... well, as anything. The Internet is not for anything in particular. It thus is for anything we invent. The Internet is not optimized for any particular use. It thus is equally optimized for anything use we might create for it.

This is its genius.

The Net Neutrality argument in the United States

But, we want to know what the Internet is. Does anything follow from the fact that it has a particular architecture?

Perhaps. If this is the Internet's architecture, then if we decide that that architecture is worth preserving, certain policy decisions are called for. This indeed helps explain one of the largest controversies in the United States over the shape of the Internet.

The Internet in the United States is delivered by a small handful of very large companies. Competition among these is further muted by the fact that generally only one or two companies provide service in any particular area. There are good business reasons for this: The cost of laying wires or cable is quite high. Further, putting in the wires requires permission from the local government which often is reluctant to allow a second, third, or fourth company to hang redundant wires or to dig up the streets yet again.

The companies that provide Internet access in the U.S. generally are not primarily in the Internet access business. They are phone companies or cable television companies that are structured to make money by selling content and services over the wires they have attached to residences. The telephone companies charge for telephone calls and for additional services such as call waiting. The cable companies charge for access to television channels and premium content. These business models are quite different from models that would encourage maximum use of the Internet and maximum innovation on the Internet.

For example, the cable company Comcast decided it would throttle the amount of traffic its users were generating using the BitTorrent peer-to-peer protocol.⁴ BitTorrent is a highly efficient way of downloading large files. It is used by many people sharing copyrighted materials, but it's also used by governments, media companies, and others. Likewise, Verizon a few years ago blocked instant messages supporting abortion rights

⁴ Soghoian, Chris. "Is Comcast's BitTorrent filtering violating the law?" CNet News. Sept. 4, 2007. http://news.cnet.com/8301-13739_3-9769645-46.html

in the United States.⁵ While the motives for doing so may be understandable, both of these actions violate the basic clause of the Internet protocol: Networks agree to pass all packets equally.

Opponents of these actions popularized the phrase “Net neutrality” (coined by Tim Wu, a lawyer at Columbia University⁶). Net neutrality is a policy that forbids those who provide Internet access from discriminating among packets based upon their content, origin, or application type. Net neutrality would forbid Comcast from blocking packets because they are being used by the Bittorrent protocol, and Verizon would be forbidden from blocking packets because they contain political expressions with which they disagree. The real fear driving supporters of Net neutrality is that the Internet access providers will be inevitably tempted to give preference to the content they’re selling, slowing the content they’re competing with. For example, they might decide that the movies their customers rent from them ought to be delivered faster — and with fewer “jitters” — than movies rented from competitors, or user-created movies from sites such as YouTube. Or, they might decide that their own telephone-over-the-Internet services ought to get preference over services such as Skype.

Net Neutrality can therefore be seen as a political policy that expresses the Internet’s architectural end-to-end policy, enshrining in governmental policy the idea that the Internet is not for anything in particular. All packets should flow equally not only because free speech is good, but because non-discrimination among packets means that new ideas can compete effectively with incumbent ideas.

Opponents of Net neutrality say that they need to discriminate among packets in order to manage their network responsibly.⁷ Suppose a user spends all day downloading large

⁵ Liptak, Adam. “Verizon Blocks Message of Abortion Rights Group.” New York Times. Sept. 27, 2007. <http://www.nytimes.com/2007/09/27/us/27verizon.html>

⁶ “Professor Tim Wu Praises FCC Plan on Net Neutrality”, Columbia Law School, Sept. 21, 2009. http://www.law.columbia.edu/media_inquiries/news_events/2009/september2009/netneutrality

⁷ For example, Larry Downes’ “The case against the FCC’s Net neutrality plan,” CNet News, Oct. 29, 2009. http://news.cnet.com/8301-1035_3-10385865-94.html, and Richard Bennett’s “Designed for Change” <http://www.itif.org/files/2009-designed-for-change.pdf>

files that infringe copyright. Why should that person's use of the Internet be allowed to degrade the use of others sharing the same line? Or, say opponents of Net neutrality, they need to be able to block the tidal wave of packets that flood the network during "denial of service" (DOS) attacks by hackers. Proponents of Net neutrality respond that there are often other alternatives available to manage the network, but if there aren't (for example, during a DOS attack), some minimal and temporary forms of discrimination are permissible.

The United States Federal Communications Commission has come out strongly in favor of Net neutrality, allowing for exceptions to keep networks up and running, as has President Obama himself. It is not exactly clear yet how it will be implemented or what the range of exemptions will include, but for now, Net neutrality is the official policy of the United States. The architecture of the Net has been taken as a value and has been written into policy.

Does the Internet have values?

That's at the policy level. And the policy only reflects the architecture of the Internet because a political process decided to take that architecture as a value worth protecting: Net neutrality implies that we think the end-to-end architecture of the Net has a value worth preserving. Are there values embedded in the Internet or only added to it?

The creators and designers of the Internet had to have been guided by some values because otherwise they wouldn't have created it. In that sense, yes the Internet expresses values. They set out to create a reliable way of moving bits, at low cost, and in a network that can expand easily by hooking in new networks and users. There is clearly embedded in the Internet a belief that communicating information has value.

But, the creators seemed to have in mind other ideas and values that make the open communication of information seem like a good thing. The Internet's architecture expresses values by what that architecture leaves out:

- The Internet is not structured to enable easy control of the content that moves over it. It thus favors free speech.
- The Internet is not structured to make it easy to associate content with its creators. It thus favors anonymity.
- The Internet does not make it easy to keep content from being shared. It thus favors sharing.
- The Internet does not make it easy to distributed content under controlled circumstances. It thus favors non-monetary transactions.

Further, behind these architectural decisions there seems to be a set of value-based beliefs about what we will do with an open network, beliefs that express ideas about human nature:

- Humans are inventive, and thus will figure out ways to put an open network to use
- Human creativity is best enabled by not constraining it
- Left to themselves, humans will organize themselves without coercion and not fall into mere chaos

This is a fairly optimistic view of humanity. And, I don't want to overstate the case. These values can all be overridden. We know for a fact that people, organizations and nations can censor content, charge for it, and track down those who refuse to obey the rules. Nevertheless, the Internet's architecture does reflect values. If the creators of the Internet had thought that it was highly valuable to make sure that the only information distributed had been validated or authenticated, they would have designed the Internet differently. They would have built it differently if they had thought it crucial that creators be assured of payment for their creations. In fact, Ted Nelson's Xanadu project⁸ had many of the elements of the Internet and the Web, except that its architecture was far

⁸ <http://www.xanadu.net>. See also "Project Xanadu," Wikipedia http://en.wikipedia.org/wiki/Project_Xanadu

more centralized because Nelson thought it essential to the health of the network that links be bidirectional (if A links to B, then B links to A) and that creators be assured that they could be paid for every use of their content. The Internet is structured differently from Xanadu because the creators of the Internet had a different idea about the network's value -- and a different view of human motivation.

The Internet thus starts off with some values expressed by what it makes easy to do and what it makes hard to do. But the openness of the Internet -- the fact that it is not designed for any one purpose and is equally open to anyone who wants to create a new service using its facilities -- makes it uniquely possible to be turned into something that spurns the values that created it. Lawrence Lessig in 1999's *Code and Other Laws of Cyberspace*⁹ points to four forces that can regulate the Internet: law, norms, markets, and architecture. Through various combinations, the Internet can become a closed network -- closed to types of information, to types of people, to particular ideas, to particular services and uses.

This makes it uniquely difficult to say what the Internet is, especially when looked at across multiple cultures. In one culture it might be the open source of all ideas. In another it might be the source of pernicious foreign influence, and thus might be filtered with some degree of success. In another culture, the Internet might be a way to improve government services but also a source of internal instability. In each case, structures can be built to change the Net's openness into something else, whether it's Italy's requirement that sessions at Internet cafes not be anonymous¹⁰ to China's insistence that users need a license (complete with authenticated photo identification) to create a new Web site.¹¹

It's getting harder to say what the Internet is.

⁹ Lessig, Lawrence. *Code and Other Laws of Cyberspace* (Basic Books, NY: 1999)
<http://www.code-is-law.org/>

¹⁰ Celeste, Sofia. "Want to check your e-mail in Italy? Bring your passport." *Christian Science Monitor*. Oct. 4, 2005. <http://www.csmonitor.com/2005/1004/p07s01-woeu.html>

¹¹ Dong, Donnie. "Wanna set up a personal website in China?" Feb. 23, 2010
<http://english.blawgdog.com/2010/02/wanna-setup-personal-website-in-china.html>

The Internet age gap(s)

And it's getting harder still. Even within one culture — the United States, for example — the Internet can be different things to different people.

My colleagues Jon Palfrey and Urs Gasser study “digital natives” — those who were born after 1980 and thus grew up with the Internet available to them — trying to understand their attitudes toward broad issues such as privacy, learning styles, and sense of self. In their book, *Born Digital*¹², Palfrey and Gasser say, for example, that digital natives often learn by skimming works, finding an area into which they dive deeply, and then synthesize their ideas by engaging in public or semi-public discussion with their social circle (pp. 240-243). Digital natives also avoid the hard work of writing papers by copying and pasting from online sources. Similarly, S. Shyam Sundar, in an essay in a volume from the MacArthur Foundation¹³, finds that digital natives often rely on superficial cues as markers of credibility, including the design and layout of the source (p. 76).

These are good researchers, so let's assume their research is accurate. The Internet therefore looks very different to digital natives than it does to the generation that created the Internet. And not just the Internet is different for the two groups. The very idea of research, learning, education, and knowledge are different because these groups treat the Internet differently.

This seems to make it less likely than ever that we can say what the Internet is. To one group it may be a swamp of misinformation, while to another it is the main source of reliable information their world.

¹² Palfrey, Jon; Gasser, Urs. *Born Digital*. (Basic Books, NY 2008)

¹³ S. Shyam Sundar. “The MAIN Model: A Heuristic Approach to Understanding Technology Effects on Credibility.” *Digital Media, Youth, and Credibility*. pp. 73-100. Miriam, Metzger and Andrew Flanagin, eds. (MIT Press, Cambridge, 2008)

But the problem is worse than that, for as researchers such as danah boyd have shown¹⁴, even within a single age group, the Internet means different things based on economic class and race. Boyd shows that even the simple choice of which social networking site to join — Facebook or Myspace —has racial and socio-economic dimensions. A recent survey by the respected group Pew Research Center¹⁵ supports boyd’s research:

White Millennials [*what Palfrey and Gasser call “digital natives”*] are the most likely to have created a social networking profile (83%). By comparison, 71% of blacks and 52% of Hispanics have done that. But among those who have created their own profile, blacks are more likely to use these sites multiple times a day (45% vs. 25% of whites).

Millennials who have attended college are significantly more likely than those with less educational attainment to have their own social networking profile; 86% of those with at least some college experience have created their own social networking profile, compared with 59% of those with no college experience. Similarly, 60% of social networking users who have attended college visit these sites at least once day; of those who have not attended college, fewer visit the sites daily (44%). (p. 29)

What is the Internet? It seems to depend on culture, age, race, class, ethnicity, educational level ... on just about *everything*.

Technodeterminism, its opposite, and its merging

So far I have been making the case that although the Internet has some values that are expressed in its architecture, what the Internet is to us varies widely. I am thus making the case against “technodeterminism,” that is, the idea that a technology has an effect all by itself and independent of culture and other such variables.

¹⁴ boyd, danah. “Viewing American class divisions through Facebook and MySpace” June 24, 2007 <http://www.danah.org/papers/essays/ClassDivisions.html>

¹⁵ *Millennials: A Portrait of Generation Next*. Pew Research Center. February, 2010.

Technodeterminism is a term that is used almost always by those who disagree with it; it is exceedingly hard to find someone who willingly calls herself a technodeterminist. Yet, some of us, including me, believe that the Internet holds transformative power. How can a technology have transformative power if what it is depends on how we take it?

The anti-technodeterminists are of course correct. Technology does nothing by itself. Further, believing in technodeterminism can lead to foolish, wasteful programs. For example, if technodeterminism were true in its strong sense, then we could just drop Internet-enabled devices -- laptops or mobile phones, perhaps -- on poor parts of the world, and the Internet would work its magic. The One Laptop Per Child project (which I believe does much good work) sometimes sounds as if that's what they believe: Give a child a laptop, and get out of the way.¹⁶ The anti-technodeterminists respond that that invests the laptop with magical power, when in fact it is merely a thing that has meaning, utility, and purpose only within a cultural context.

I believe both the technodeterminists and the anti-technodeterminists. The Internet by itself possesses no magic. It can be used as a tool of repression as well as a tool of freedom. It can be manifest within a society as a hope for the advancement of one's culture or as a threat to its integrity. And yet...

...And yet, there seem to me to be certain basic lessons that anyone, in any global culture, learns simply by being on the Internet and browsing the Web. You learn:

- There are many people with many ideas
- Not everyone agrees ... about anything
- It is possible to post an idea to the world without being a major media power
- Many people will contribute their ideas and creativity for free

¹⁶ The founders of the OLPC program favor a version of Constructionism, based on Constructivism, that stresses peer-to-peer learning, and not "being force fed information." See the Constructionist article on the OLPC's web site: <http://wiki.laptop.org/go/Constructionist>. See also Ethan Zuckerman's "One Laptop Per Child: Just what sort of content do you load onto these puppies?" <http://www.ethanzuckerman.com/blog/2007/01/22/one-laptop-per-child-just-what-sort-of-content-do-you-load-onto-these-puppies/>

- Ideas can be linked
- Ideas can lead in many directions
- There is almost always more information to be gotten
- There is too much to know, to explore
- Around the world, people are speaking from their hearts
- We are all interested in what happens to us
- The world is more interesting than we could have imagined
- We have responsibility for making sense of our world and our lives

It seems to me that those are inevitable lessons we learn just by being on the Internet.

On the other hand, what we make of those lessons is highly dependent on our culture, age, gender, and so forth: We may see all the disagreements on the Net and react in horror. We may think that the Net is therefore a source of overwhelming falsehood and dissatisfaction. Or we may think that the Net is therefore a place of joy and delight. We may think whatever our particular circumstances lead us to think. The Internet does not determine us to take its nature any particular way. But it does have a nature that we have to take one way or another.

What is the Internet? A noisy peace.

So, what is the Internet?

It has an architecture that, perhaps unique among technologies, is designed to allow it to be used for unanticipated purposes. It thus can't be defined by the services it provides, because it is more than the sum of those services. It is the possibility of building those services and the ones we haven't imagined yet.

But, as that technology -- defined by its openness -- is taken up by various cultures, its openness looks different. That's inevitable since the differences among cultures about openness are among the most intimate.

So, let me tell you how the Internet looks to me, and why I am so hopeful about it. But, I recognize that this view is conditioned by the fact that I am of a particular setting: I am American, 59 years old, white, male, Jewish, intellectual, healthy, middle class ... All of those (and more) condition how I think. I don't believe I can step outside of those factors. I do believe it is important to acknowledge them, especially given what I am about to say.

In the West, we have generally assumed that to know something is to have driven out disagreement; if reasonable people still disagree, we say "We don't know that yet." We have generally supported the Enlightenment ideal that if two reasonable people who disagree sit down and have a reasonable conversation for a long enough time, they will come to agreement.

I think the Internet disproves that. One of the lessons of the Internet is, I believe, that we human beings are never going to agree about anything. No fact goes uncontradicted on the Net. No belief escapes contravention. No value is not flaunted. In this, the Internet is merely revealing the state of the world. Back when what we knew of the world came filtered through controlled media — books, radio, television — it was easier to avoid this basic fact. Now we can see all those contradictions. Now we cannot avoid those contradictions.

So, the Internet places a new challenge in front of us: Can we forge a peace among us that does not rest upon everyone agreeing, but upon all of us disagreeing in peaceful ways?

And I'd like to suggest that the Internet not only forces us to acknowledge that we humans disagree about everything. It gives us a new tool for disagreeing peaceably: the hyperlink. Hyperlinks enable us not only to hear distant voices, but to link to them. The disagreements are thus held close to us, through links.

Now, this is a slim hope. If we are to form peace out of perpetual disagreement, the links will have to come bottom up, from citizens who are listening to others, linking to them, and explaining those differences. This is something any citizen with an Internet connection can do. My friend and mentor Ethan Zuckerman reminds us of the particular importance of “bridge figures,”¹⁷ that is, people who are able to translate from one culture to another — translate not just language, but cultures, values, and points of view. If the Internet lets us hear all those differing voices, hyperlinks enable us to engage with them, begin to make sense of them, and perhaps learn to live with them.

But, for this hope to be real -- and of course hyperlinks can also be used to whip up furious hatred — then we need the Internet to be open: open to every dweller on the planet, open to every idea, open to every disagreement. Such an Internet would reflect the Internet’s technical architecture. But, as we have seen, it does not have to do so. We will have such an Internet only if we choose it, personally and politically. We will choose to do so only if we have hope — hope similar to what drove the Internet’s creators — that from linked difference can emerge a noisy peace. The Internet would then have a nature and a definition, but only because in an act of hopeful will, we chose to give it one.

¹⁷ For example, see “From protest to collaboration: Paul Simon’s ‘Graceland’ and lessons for xenophiles” <http://www.ethanzuckerman.com/blog/2009/04/02/from-protest-to-collaboration-paul-simons-graceland-and-lessons-for-xenophiles/>