The Theory of Generativity

David G. Post[1]

Jonathan Zittrain's "The Generative Internet"[2] comes bearing high expectations. The virtual ink on Prof. Zittrain's screen was hardly dry before Prof. Lawrence Lessig, cyberlaw's most influential scholar, gave Zittrain's argument a nickname ("Z-theory"), deemed it "brilliant," and claimed that it solved a serious problem that had been plaguing cyberlaw ever since Lessig's own *Code and Other Laws of Cyberspace* was published in 1999 – a long time, in cyber-years.[3] And Zittrain himself describes his agenda in ambitious terms; after identifying three separate, and often-conflicting, strands of theory in recent cyberlaw scholarship, he suggests that his approach – his theory of generativity – can reconcile them all.[4]

A build-up like that, and all cyberlaw scholars must (and undoubtedly will) pay close attention. As well they should; for anyone who cares about the global network and the new forms of communication and interaction it has fostered, this is vital and important scholarship, ideas that demand to be taken seriously, and to be taken up by others in the field for debate and discussion.

Like most good arguments, Zittrain's has three propositions and a program. The propositions are:

1. Generativity is what makes the global network of interconnected PCs – "the PC/Internet grid" – so remarkable and so transformative a technology.

2. Generativity is not an inherent, or immutable, feature of the PC/Internet grid; it is, rather, the result of a number of specific and identifiable design decisions implemented in code(s) – primarily, PC operating systems and the TCP/IP protocols; these codes can be re-designed, re-engineered, and re-written to make the grid a much less generative place.

3. Security vulnerabilities – traceable ultimately to the very features of the grid that make it generative – will be the crucial lever in causing consumers and government regulators alike to demand a safer and more secure (albeit much less generative) environment, which hardware and software manufacturers and service providers will be only to happy to provide.

---

[1] I. Herman Stern Professor of Law, Beasley School of Law, Temple University. David.Post@temple.edu; http://www.davidpost.com.

[2] Jonathan L. Zittrain, "The Generative Internet," 119 Harv. L. Rev 1974 (2006) (hereinafter, "The Generative Internet"). My comments below presuppose a reader who is already familiar with Zittrain's article; my attempts to summarize Zittrain's arguments are necessarily sketchy and incomplete, and I would not want them to dissuade anyone from tackling the original.

[3] Lawrence Lessig, *Code version 2.0* (Basic Books, 2006), at 74-76.

[4] The Generative Internet, at 2002 (describing "the three formerly competing strands of cyberlaw scholarship" that can now be "reconciled to explain the path of the Internet")

The program? Minimize the damage that will, in the face of this intense pressure from all sides, inevitably occur; try to establish "principles that will blunt the most unappealing features of a more locked-down technological future"[5]; help identify those modifications to the PC/Internet grid that will "do the least harm to its generative possibilities."[6]

1. Generativity is what makes the global network of interconnected PCs – the "PC/Internet grid" – so remarkable and so transformative a technology.

> "Generativity denotes a technology's overall capacity to produce unprompted change driven by large, varied, and uncoordinated audiences. The grid of PCs connected by the Internet has developed in such a way that it is consummately generative. From the beginning, the PC has been designed to run almost any program created by the manufacturer, the user, or a remote third party and to make the creation of such programs a relatively easy task. When these highly adaptable machines are connected to a network with little centralized control, the result is a grid that is nearly completely open to the creation and rapid distribution of the innovations of technology-savvy users to a mass audience that can enjoy those innovations without having to know how they work."[7]

Open machines, capable of executing code produced by any third parties, all connected to the open network, the one "that no one in particular owns and that anyone can join."[8] Code can come from anywhere on the grid, it can be instantaneously distributed to every point on the grid, and it can be implemented and used by anyone connected to the grid – a recipe for an innovation-generating machine the likes of which we've never quite seen before. It is, as Zittrain puts it, "difficult to identify in 2006 a technology bundle more generative than the PC and the Internet to which it attaches."[9]

It's pretty obvious, when you think about it this way – and I mean that as a compliment; "stating the obvious," articulating a proposition in a manner that (pretty much) everyone involved in the discussion agrees *is* obvious, can move discussion and debate forward in constructive ways, and Zittrain's contribution here will undoubtedly do just that for cyberlaw. Those of you who are not deeply embedded in the world of cyberlaw scholarship might find it odd that, as Zittrain puts it, the field "has struggled to identify precisely what is so valuable about today's Internet"[10]; you might have thought that those of us who spend so much of our time thinking about the Internet would, surely, have sorted that out, by now. Reading "The Generative Internet" helped me to realize that we actually *have* sorted it (or, at least, some of it) out, that there is, in fact, a kind of

---

[5] *Id.*, at 1977.78.
[6] *Id.*, at 2031.
[7] *Id.*, at 1980.
[8] *Id.*, at 1993.
[9] *Id.*, at 1982.
[10] *Id.*, at 1978.

rough consensus in the field on the question (though it has been difficult to articulate precisely where that consensus lay).

Web browsers; search engines; online auctions; peer-to-peer filesharing; Instant Messaging; iTunes; Amazon; Skype; Wikipedia; Flickr; Linux; Facebook; Slashdot; VOIP; Youtube; blogs; massively multi-player online games; virtual worlds. Roll back the clock a mere 15 years, and not a single one of those words or phrases would have made the slightest sense to anyone reading this article. That's a lot of new words, and a lot of innovation, in 15 years, innovation on top of innovation on top of innovation, carried around the globe over the consummately open network (so that it can spur others to even more innovation). It's a wonderful story, and Zittrain's article will, one hopes, make more people appreciate it. Something important, and quite possibly transformative, is going on here; cyberlaw scholars do all (more-or-less) agree about that – why else would we all be spending so much of our time writing about the Internet and its law? And "generativity" comes close, I think, to capturing what that something might be, giving us a nice, shorthand label for a crucial and complicated thought. I think that most cyberlaw scholars, whatever else they may think about the Internet and the legal problems associated with online conduct, and whatever their position on the many quite fundamental issues on which we often disagree – the proper scope of Internet regulation, the jurisdictional rules governing border-crossing activities, the nature of (and best way to safeguard) free expression in Internet communications, the proper scope of intellectual property protection, and so on – recognize, and appreciate, and even celebrate, the network's astonishing ability to generate innovative uses of itself.

2. *Generativity is not an inherent, or immutable, feature of the PC/Internet grid; it is, rather, the result of a number of specific and identifiable design decisions implemented in code(s) – primarily, PC operating systems and the TCP/IP protocols; these codes can be re-designed, re-engineered, and re-written to make the grid a much less generative place.*

> "Though the openness of the Internet and the PC to third-party contribution is now so persistent and universal as to seem inevitable, these technologies need not have been [so] configured."[11]

Right, again. Generativity is "designed into" the grid; it is a consequence (at least in significant measure) of the Internet's basic "end-to-end" architecture, and the similarly open architecture – "open" in the sense of allowing any third-party applications, and third-party devices, to inter-connect and inter-operate – of the PC. End-to-end design mandates that the network be kept as simple and stupid as possible; the network does nothing more than move bits from one point to another, with all of the processing of those bits – the message-reading, and message-interpreting, and message-checking stuff – taking place at the end-points of the network. The Internet has

> ". . . an hourglass design, with a simple set of narrow and slow-changing protocols in the middle, resting on an open stable of physical carriers at

[11] *Id.*, at 1976-77.

> the bottom and any number of applications written by third parties on the top."[12]

This network doesn't care what you're doing "below" (*i.e.*., what kind of physical hardware you use to send and receive bits) or "above" (*i.e.*, what kind of software you use to make sense of those bits that you send and receive); it stands in the middle of the hourglass, taking whatever bits your hardware delivers to it, and moving those bits to wherever it is directed to move them, and leaving all tasks of interpretation and processing to the recipient.

PCs, also, have this kind of hourglass design:

> "[A] similar principle operates to push processing to the third-party applications running on, and not the operating system running in, those machines. . . . The Internet hourglass, despite having been conceived by an utterly different group of architects from those who designed or structured the market for PCs, . . . mirrors PC architecture in key respects. The network is indifferent to both the physical media on which it operates and the nature of the data it passes, just as a PC is open to running upon a variety of physical hardware 'below' and to supporting any number of applications from various sources 'above.'"[13]

This is not, Zittrain reminds us, inevitable. There are lots of other ways to design machines that can connect to one another over a network, and lots of other ways to design networks to connect those machines together. We have many examples of such machines and such networks all around us (and many more if we want to spend some time digging around in the technology graveyards). For example, the machines we call "telephones," and "televisions," and "radios," and "fax machines," and "TiVos," (and those we used to call "telegraphs," and "telexes," . . .) – all are capable of being networked together, and all are capable of doing some wonderful things, but all have "operating systems" that are most emphatically not open to third-party code. Similarly, many networks aren't end-to-end in design, but instead put most of their "intelligence" and processing power *inside* the network (rather than at the network end-points). The POTS (Plain Old Telephone Service) network is only the most prominent of many such examples: a very smart network, relying on closed network protocols to connect together very dumb machines at the end-points of the network in a very un-generative fashion, the very opposite of the end-to-end Internet grid (smart end-points, dumb network).[14]

---

[12] *Id.*, at 1988.

[13] *Id.*, at 1988-89.

[14] *See, e.g.*, National Research Council, *The Internet's Coming of Age* (Nat'l Acad. Press 2001), at 81-85 (describing telephone network structure). It's not just the telephone network that is configured in this ungenerative manner; many computer networking and inter-networking protocols that were alternatives to TCP/IP in the 1980s and 1990s – *e.g.*, the International Telecommunications Union's X.25 and OSI protocols, Digital Equipment's DECNET, IBM's System Network Architecture, among many others – are not designed along end-to-end lines. *See, generally*, National Research Council, *Global Networks and Local Values* (Nat'l Acad. Press, 2001), at 23-45; Janet Abbate, *Inventing the Internet* (MIT Press, 2000), at 147-179.

But with all these different kinds of machines, and all these different ways of connecting them together, *the Big One turned out to be the one that connects the generative machines (PCs) to the generative network (the TCP/IP network).* Duh! It's so obvious! Hook the generative machines up to the generative network and it will – well, it will *generate*. It will swamp all of the others precisely because it is so generative. The global telephone and television (and fax, and telex, and telegraph, etc.) networks were (and in some cases are still) wonderful things, in many ways; but the kind of innovation they stimulate was, and is, of a different, and much lower, order than that of the PC/Internet grid. They might survive (though personally I doubt it); but they surely will never again dominate, because they can't keep up; they can't grow (as fast); they don't – literally – generate.[15]

Now, all of this, Zittrain reminds us, can change. It's all written in code, not stone, implemented in PC and network hardware and software in any number of identifiable places, and code can be changed. PC operating systems don't have to allow execution of third-party code; that is not, somehow, part of their immutable "nature."[16] Nor do the Internet protocols have to call for the transmission of any and all bit-streams, from any and all sources, to any and all recipients. All of this can change; PCs can become "appliancized" – more like TiVo boxes. And the inter-network can be made more "intelligent" than it is now; network routers can be programmed to do a lot more processing than they do now – to check for transmission errors, or for "spam," or for hostile code (viruses and the like), or for "unauthorized" users, to "authenticate" messages on the basis of user identity, or geographic location, to "discriminate" among messages based on their content or the identity of the communicating parties, etc.

And these changes, to be sure, would likely make this a much, much less generative place – a bad thing, one would think (see proposition number 1).

---

[15] The parallel to those other "information systems" we classify as "biological" is, incidentally, quite precise: the most generative ones prevail, because they are the most generative ones.

[16] Zittrain's argument here is part of a larger "nature vs. nurture" debate in cyberlaw – another interesting facet of the scholarship that might be unfamiliar to outsiders. Much of Lawrence Lessig's influential *Code and Other Laws of Cyberspace* was devoted to showing that the Internet has no "nature," no immutable or inherent characteristics (and, in particular, that it is not "by nature" or "inherently" unregulable); we built it to be the way it is, Lessig argued, and we can build it differently if we want to.

The semantics of the debate still bother me. Lessig (and Zittrain) are surely correct about mutability; the Internet is a built environment, and its design can be altered (to make it more regulable, to abandon end-to-end features, etc.). It does not follow, though, that because it is mutable, it has no "nature." To describe the "nature" of things doesn't imply immutability – whether you're talking about the "nature" of human intelligence, the "nature" of the HIV virus, the "nature" of a tropical forest, or the "nature" of the Internet; it is surely in the "nature" of a tree to grow towards the sunlight, and just as surely we can rather easily prevent that from happening (*i.e.*, by chopping it down), just as it is part of many peoples' "nature" – embedded, in fact, in their genetic code – to produce an insufficient quantity of insulin, another highly mutable characteristic. Semantic confusion on this score has a long pedigree, and though it is just semantics, it can have unfortunate implications for the ongoing discussion; in cyberlaw, in my opinion, it has had the unfortunate effect of discouraging people from asking useful and important questions about the "nature" of the Internet.

But who would want that to happen? Who would want to kill the goose that is laying the golden eggs?

3. *Security vulnerabilities – traceable ultimately to the very features of the grid that make it generative – will be the crucial lever in causing consumers and government regulators alike to demand a safer and more secure (albeit much less generative) environment, which hardware and software manufacturers and service providers will be only to happy to provide.*

Now things get really interesting.

> "[T]he slope of this innovative curve may . . . soon be constrained by some of the very factors that have made it so steep. Such constraints may arise because generativity is vulnerability in the current order: the fact that tens of millions of machines are connected to networks that can convey reprogramming in a matter of seconds means that those computers stand exposed to near-instantaneous change. This kind of generativity keeps publishers vulnerable to the latest tools of intellectual property infringement, crafted ever more cleverly to evade monitoring and control, and available for installation within moments everywhere. It also opens PCs to the prospect of mass infection by a computer virus that exploits either user ignorance or a security vulnerability that allows code from the network to run on a PC without approval by its owner."[17]

Zittrain's point here is not the obvious one (and, in this case, the uninteresting one) that people can use open PCs, connected to an open network, for bad things as well as for good things, that these tools, like all tools, can be used for good or ill. It is, rather, that generativity is itself both the good *and* the ill, at war with itself, bearing within itself the seeds of its own destruction. [Fans of Wagner's "Ring" cycle will particularly appreciate Zittrain's argument here.] Generativity comes from openness, and the more open the system, the more vulnerable it is to attack. Open machines on an open network are (literally) an open invitation – to the virus-writers, the spammers, the phish-ers, the denial-of-service attackers, and all the rest.[18] As this remarkable capability we have developed – open machines attached to the open network – starts to look less and less appealing to more and more people, it will "precipitate an unprecedented intervention into the way today's Internet and PCs work,"[19] a backlash that "makes it not merely

---

[17] "The Generative Internet," at 1995 (emphasis in original)

[18] "Combine one well-written worm of the sort that can evade firewalls and antivirus software with one truly malicious worm writer, and worrisome low-level annoyances could spike to more acute effects: checkin unavailable at some airline counters; no overnight deliveries or other forms of package and letter distribution; payroll software unable to generate paychecks for millions of workers; or vital records held in medical offices, schools, town halls, and other data repositories eliminated, released, or nefariously altered." The Generative Internet, at 2012-13.

[19] *Id*., at 1977.

possible, but likely, that the personal computing and networking environment of tomorrow will be sadly hobbled, bearing little resemblance to the one that most of the world enjoys today."[20]

Many willing and eager participants would "appreciate a world in which the PC is more locked down,"[21] and they will lead the backlash: Consumers, fed up with machines constantly invaded by, or infected with, or swamped under garbage (or worse),[22] government regulators who "would welcome and even encourage a PC/Internet grid that is less exceptional and more regulable,"[23] and "mature technology industry players" – hardware and software providers happy to give consumers what they want (a safer Internet experience) and happy, as yesterday's innovators often are, to make things more difficult for tomorrow's innovators. The result? The spread of "Internet-aware appliances that exploit the generativity of OSs and networks but that are not themselves generative,"[24] and the "appliancization" of the PC itself.

> "The most plausible path along which the Internet might develop is one that finds greater stability by imposing greater constraint on, if not outright elimination of, the capacity of upstart innovators to demonstrate and deploy their genius to large audiences. Financial transactions over such an Internet will be more trustworthy, but the range of its users' business models will be narrow. This Internet's attached consumer hardware will be easier to understand and will be easier to use for purposes that the hardware manufacturers preconceive, but it will be off limits to amateur tinkering."[25]

Hence, the program: We – all those who care about openness and innovation, who "treasure the Internet's generative features"[26] – need to start taking these security vulnerabilities seriously. A "free and open Internet . . . is simply not possible unless the most pressing demands of countervailing regulatory forces are satisfied"[27]; our task, Zittrain says, is to insist that "modifications to the PC/Internet grid be made [in ways that] will do the least harm to its generative possibilities,"[28] to help develop "thoughtful

---

[20] *Id.*
[21] *Id.*, at 2026.

[22] Consumers, Zittrain writes, who are "deciding between security-flawed generative PCs and safer but more limited information appliances (or appliancized PCs)," may "consistently undervalue the benefits of future innovation (and therefore of generative PCs)," because "[t]he benefits of future innovation are difficult to perceive in present-value terms, and few consumers are likely to factor into their purchasing decisions the history of unexpected information technology innovation that promises so much more just around the corner." *Id.*, at 2021.

[23] *Id.*, at 2002.
[24] *Id.*, at 2015.
[25] *Id.*
[26] *Id.*
[27] *Id.*, at 2028.
[28] *Id.*, at 2031.

interventions in law and code"[29] that will "blunt the most unappealing features of a more locked-down technological future,"[30] interventions through which the grid can be made more secure "*without* sacrificing its essential generative characteristics."[31] The stakes are high; if we fail, "consumer sentiment and preexisting regulatory pressures [may] prematurely and tragically terminate the grand experiment that is the Internet today."[32]

* * * * *

It would, perhaps, be uncharitable to point out (as Zittrain himself candidly acknowledges[33]) that there have been many prior reports of the Internet's imminent death, and that, to paraphrase Mark Twain, they have all been greatly exaggerated. Nothing is harder to predict, as they say, than the future. I've been skeptical of these various doomsday predictions in the past,[34] and I remain (somewhat) so still. Many parts of Zittrain's account of things do strike me as persuasive; user demand for a more secure and virus/spam/malware-free communications environment will undoubtedly intensify, and regulators – whether out of genuine desire to serve their constituents' interests, or as a pretext to serve their own – will surely bring more intense pressure to bear, in the name of stability and security, to rein things in and to lock things down. But it is much less clear to me than to Zittrain that the "SecureNet" of the future that he sketches out necessarily supplants the generative grid, rather than co-existing with it; a grid of 400 million open PCs is not less generative than a grid of 400 million open PCs and 500 million locked-down TiVos. I also remain skeptical, for reasons I've set forth elsewhere, that the regulators will have as straightforward a time as Zittrain appears to think they will have bringing generativity to heel; the grid may be even more generative than Zittrain allows, generating both its own techniques for avoiding the more "unfortunate [and] hamhanded" attempts to bring it under control, and its own solutions to the authentication/security dilemmas that Zittrain so properly focuses on.

But this is not only uncharitable, it is unproductive. Time will tell the extent to which his predictions pan out –itself a function, in part, of what we, collectively, do about the problem. Zittrain's program is a worthy one. Generativity matters. We need to identify where it is under attack, and we need to blunt those attacks wherever possible. We need to identify – and perhaps even to help invent – the mechanisms and processes and institutions that simultaneously address legitimate security and vulnerability concerns

---

[29] *Id.*, at 2028-29.
[30] *Id.*, at 1997.
[31] *Id.*, at 2026 (emphasis supplied).
[32] *Id.*, at 1979.
[33] *See, e.g., id.*, at 1997 (noting that while "the notion that code is law undergirds a powerful theory . . . some of its most troubling empirical predictions about the use of code to restrict individual freedom have not yet come to pass"); *id.*, at 1999 (noting that the "central fears of the 'code is law' theory as applied to the Internet and PCs . . . have largely remained hypothetical"); *id.*, at 1979-80 (no "discernable movement towards the locked-down dystopias" predicted by DRM opponents).
[34] *See, e.g.*, David G. Post, "What Larry Doesn't Get: Code, Law, and Liberty in Cyberspace," 52 Stan. L. Rev 1439 (2000).

while preserving, and perhaps even enhancing, generativity.[35] Some of us may be less willing than he to make the "concessions" he thinks are required,[36] or to abandon the "bright line against nearly any form of increased Internet regulability" that he thinks is "no longer tenable."[37] But there will be time enough to fight those battles down the road.

[35] *See, e.g.*, David Johnson, Susan Crawford, and John Palfrey, "The Accountable Internet," 9 Va. J. L. and Tech. 1 (2004). Zittrain himself identifies a number of projects that deserve support, whatever one thinks of the likelihood that his gloomy predictions will come true, *e.g.*, setting up "a technical architecture to label applications and fragments of executable code, coupled with an organization to apply such labels nondiscriminatorily," The Generative Internet, at 2033, helping to "establish a distributed architecture by which the individual decisions about whether to run a given application, and the subsequent results, could serve as advice to others contemplating whether to run such code," *id.*, or helping to develop tools to "provide members of the general Internet public with simple but powerful information about the code they encounter [such as] a dashboard displaying information such as how many other computers in the world were running a candidate piece of software and whether their users were on average more or less satisfied with their computers than those who did not run it." *Id.*, at 2034.

[36] *Id.*, at 2028.
[37] *Id.*, at 1979.