CODE

version 2.0

LAWRENCE LESSIG

BASIC BOOKS

A Member of the Perseus Books Group
New York

FIVE

regulating code

COMMERCE HAS DONE ITS PART—FOR COMMERCE, AND INDIRECTLY, FOR governments. Technologies that make commerce more efficient are also technologies that make regulation simpler. The one supports the other. There are a host of technologies now that make it easier to know who someone is on the Net, what they're doing, and where they're doing it. These technologies were built to make business work better. They make life on the Internet safer. But the by-product of these technologies is to make the Net more regulable.

More regulable. Not perfectly regulable. These tools alone do a great deal. As Joel Reidenberg notes, they are already leading courts to recognize how behavior on the Net can be reached—and regulated.[1] But they don't yet create the incentives to build regulability into the heart of the Net. That final step will require action by the government.[2]

When I wrote the first version of this book, I certainly expected that the government would eventually take these steps. Events since 1999—including the birth of Z-theory described below—have only increased my confidence. In the United States, the identification of "an enemy"—terrorism—has weakened the resolve to resist government action to make government more powerful and regulation more effective. There's a limit, or at least I hope there is, but there is also no doubt that the line has been moved. And in any case, there is not much more that the government would need to do in order to radically increase the regulability of the net. These steps would not themselves excite any significant resistance. The government has the means, and the motive. This chapter maps the opportunity.

The trick is obvious once it is seen. It may well be difficult for the government to regulate behavior directly, given the architecture of the Internet as it is. But that doesn't mean it is difficult for the government to regulate the architecture of the Internet as it is. The trick, then, is for the government to take steps that induce the development of an architecture that makes behavior more regulable.

In this context, I don't mean by "architecture" the regulation of TCP/IP itself. Instead, I simply mean regulation that changes the effective constraints of the architecture of the Internet, by altering the code at any layer within that space. If technologies of identification are lacking, then regulating the architecture in this sense means steps the government can take to induce the deployment of technologies of identification.

If the government takes these steps, it will increase the regulability of behavior on the Internet. And depending upon the substance of these steps taken, it could render the Internet the most perfectly regulable space we've known. As Michael Geist describes it, "governments may have been willing to step aside during the commercial Internet's nascent years, but no longer."[3]

REGULATING ARCHITECTURE: THE REGULATORY TWO-STEP

We can call this the "regulatory two-step": In a context in which behavior is relatively unregulable, the government takes steps to increase regulability. And once framed, there are any number of examples that set the pattern for the two-step in cyberspace.

Car Congestion

London had a problem with traffic. There were too many cars in the central district, and there was no simple way to keep "unnecessary" cars out.

So London did three things. It first mandated a license plate that a video camera could read, and then it installed video cameras on as many public fixtures as it would take to monitor—perpetually—what cars were where.

Then, beginning in February 2003, the city imposed a congestion tax: Initially £5 per day (between 7 A.M. and 6:30 P.M.) for any car (save taxis and residents paying a special fee), raised to £8 in July 2005. After 18 months in operation, the system was working "better than expected." Traffic delays were down 32 percent, traffic within the city was down 15 percent, and delays on main routes into the zones were down 20 percent. London is now exploring new technologies to make it even easier to charge for access more accurately. These include new tagging technologies, as well as GPS and GSM technologies that would monitor the car while within London.[4]

Telephones

The architecture of telephone networks has undergone a radical shift in the past decade. After resisting the design of the Internet for many years,[5] telephone networks are now shifting from circuit-switched to packet-switched networks. As with the Internet, packets of information are spewed across the system, and nothing ensures that they will travel in the same way, or along the same path. Packets take the most efficient path, which depends on the demand at any one time.

This design, however, creates problems for law enforcement—in particular, that part of law enforcement that depends upon wiretaps to do their job. In the circuit-switched network, it was relatively simple to identify which wires to tap. In the packet-switched network, where there are no predictable paths for packets of data to travel, wiretapping becomes much more difficult.

At least it is difficult under one design of a packet-switched network. Different designs will be differently difficult. And that potential led Congress in 1994 to enact the Communications Assistance for Law Enforcement Act (CALEA). CALEA requires that networks be designed to preserve the ability of law enforcement to conduct electronic surveillance. This requirement has been negotiated in a series of "safe harbor" agreements that specify the standards networks must meet to satisfy the requirements of the law.

CALEA is a classic example of the kind of regulation that I mean this chapter to flag. The industry created one network architecture. That architecture didn't adequately serve the interests of government. The response of the government was to regulate the design of the network so it better served the government's ends. (Luckily for the networks, the government, at least initially, agreed to pick up part of the cost.[6]) As Susan Crawford writes,

> Most critically for the future of the Internet, law enforcement . . . has made clear that it wants to ensure that it reviews all possibly relevant new services for compliance with unstated information-gathering and information-forwarding requirements before these services are launched. All prudent businesses will want to run their services by law enforcement, suggests the DOJ: "Service providers would be well advised to seek guidance early, preferably well before deployment of a service, if they believe that their service is not covered by CALEA. . . . DOJ would certainly consider a service provider's failure to request such guidance in any enforcement action."[7]

CALEA is a "signal," Crawford describes, that the "FCC may take the view that permission will be needed from government authorities when designing a wide variety of services, computers, and web sites that use the Internet protocol. . . . [I]nformation flow membranes will be governmentally mandated as part of the design process for online products and services."[8] That hint has continued: In August 2005, the Federal Communications Commission (FCC) ruled that Voice-over-IP services "must be designed so as to make government wiretapping easier."[9]

Of course, regulating the architecture of the network was not the only means that Congress had. Congress could have compensated for any loss in crime prevention that resulted from the decreased ability to wiretap by increasing criminal punishments.[10] Or Congress could have increased the resources devoted to criminal investigation. Both of these changes would have altered the incentives that criminals face without using the network's potential to help track and convict criminals. But instead, Congress acted to change the architecture of the telephone networks, thus using the networks directly to change the incentives of criminals indirectly.

This is law regulating code. Its indirect effect is to improve law enforcement, and it does so by modifying code-based constraints on law enforcement.

Regulation like this works well with telephone companies. There are few companies, and the regulation is relatively easy to verify. Telephone companies are thus regulable intermediaries: Rules directed against them are likely to be enforced.

But what about when telephone service (or rather "telephone service") begins to be carried across the Internet? Vonage, or Skype, rather than Bell South? Are these entities similarly regulable?[11]

The answer is that they are, though for different reasons. Skype and Vonage, as well as many other VOIP providers, seek to maximize their value as corporations. That value comes in part from demonstrating reliably regulable behavior. Failing to comply with the rules of the United States government is not a foundation upon which to build a healthy, profitable company. That's as true for General Motors as it is for eBay.

Telephones: Part 2

Four years after Congress enacted CALEA, the FBI petitioned the Federal Communications Commission to enhance even further government's power to regulate. Among the amendments the FBI proposed was a regulation designed to require disclosure of the locations of individuals using cellular phones by requiring the phone companies to report the cell tower from which the call was served.[12] Cellular phone systems need this data to ensure seamless switching between transmitters. But beyond this and billing, the phone companies have no further need for this information.

The FBI, however, has interests beyond those of the companies. It would like that data made available whenever it has a "legitimate law enforcement reason" for requesting it. The proposed amendment to CALEA would require the cellular company to provide this information, which is a way of indirectly requiring that it write its code to make the information retrievable.[13]

The original motivation for this requirement was reasonable enough: Emergency service providers needed a simple way to determine where an emergency cellular phone call was coming from. Thus, revealing location data was necessary, at least in those cases. But the FBI was keen to extend the reach of location data beyond cases where someone was calling 911, so they pushed to require the collection of this information whenever a call is made.

So far, the FBI has been successful in its requests with the regulators but less so with courts. But the limits the courts have imposed simply require the FBI to meet a high burden of proof to get access to the data. Whatever the standard, the effect of the regulation has been to force cell phone companies to build their systems to collect and preserve a kind of data that only aids the government.

Data Retention

Computers gather data about how they're used. These data are collected in logs. The logs can be verbose or not—meaning they might gather lots of data, or little. And the more they gather, the easier it will be to trace who did what.

Governments are beginning to recognize this. And some are making sure they can take advantage of it. The United States is beginning to "mull,"[14] and the European Union has adopted, legislation to regulate "data generated or processed in connection with the provision of publicly available electronic communications," by requiring that providers retain specified data to better enable law enforcement. This includes data to determine the source, destination, time, duration, type, and equipment used in a given communication.[15] Rules such as this will build a layer of traceability into the platform of electronic communication, making it easier for governments to track individual behavior. (By contrast, in 2006, Congressman Ed Markey of Massachusetts proposed legislation to forbid certain Internet companies, primarily search engines, from keeping logs that make Internet behavior traceable.[16] We'll see how far that proposed rule gets.)

Encryption

The examples so far have involved regulations directed to code writers as a way indirectly to change behavior. But sometimes, the government is doubly indirect: Sometimes it creates market incentives as a way to change code writing, so that the code writing will indirectly change behavior. An example is the U.S. government's failed attempt to secure Clipper as the standard for encryption technology.[17]

I have already sketched the Janus-faced nature of encryption: The same technology enables both confidentiality and identification. The government is concerned with the confidentiality part. Encryption allows individuals to make their conversations or data exchanges untranslatable except by someone with a key. How untranslatable is a matter of debate,[18] but we can put that debate aside for the moment, because, regardless, it is too untranslatable for the government's liking. So the government sought to control the use of encryption technology by getting the Clipper chip accepted as a standard for encryption.

The mechanics of the Clipper chip are not easily summarized, but its aim was to encourage encryption technologies that left a back door open for the government.[19] A conversation could be encrypted so that others could not understand it, but the government would have the ability (in most cases with a court order) to decrypt the conversation using a special key.

The question for the government then was how it could spread the Clipper chip technology. At first, the Clinton administration thought that the best way was simply to ban all other encryption technology. This strategy proved very controversial, so the government then fixed on a different technique: It subsidized the development and deployment of the Clipper chip.[20]

The thinking was obvious: If the government could get industry to use Clipper by making Clipper the cheapest technology, then it could indirectly regulate the use of encryption. The market would do the regulation for the government.[21]

The subsidy plan failed. Skepticism about the quality of the code itself, and about the secrecy with which it had been developed, as well as strong opposition to any governmentally directed encryption regime (especially a U.S.-sponsored regime), led most to reject the technology. This forced the government to take another path.

That alternative is for our purposes the most interesting. For a time, some were pushing for authority to regulate authors of encryption code directly—with a requirement that they build into their code a back door through which the government could gain access.[22] While the proposals have been various, they all aim at ensuring that the government has a way to crack whatever encryption code a user selects.

Compared with other strategies—banning the use of encryption or flooding the market with an alternative encryption standard—this mode presents a number of advantages.

First, unlike banning the use of encryption, this mode of regulation does not directly interfere with the rights of use by individuals. It therefore is not vulnerable to a strong, if yet unproven constitutional claim that an individual has a right "to speak through encryption." It aims only to change the mix of encryption technologies available, not to control directly any particular use by an individual. State regulation of the writing of encryption code is just like state regulation of the design of automobiles: Individual use is not regulated. Second, unlike the technique of subsidizing one market solution, this solution allows the market to compete to provide the best encryption system, given this regulatory constraint. Finally, unlike both other solutions, this one involves the regulation of only a relatively small number of actors, since manufacturers of encryption technology are far fewer in number than users or buyers of encryption systems.

Like the other examples in this section, then, this solution is an example of the government regulating code directly so as to better regulate behavior indirectly; the government uses the architecture of the code to reach a particular substantive end. Here the end, as with digital telephony, is to ensure that the government's ability to search certain conversations is not blocked by emerging technology. And again, the government pursues that end not by regulating primary behavior but by regulating the conditions under which primary behavior happens.

REGULATING CODE TO INCREASE REGULABILITY

All five of these examples address a behavior that the government wants to regulate, but which it cannot (easily) regulate directly. In all five, the government thus regulates that behavior indirectly by directly regulating technologies that affect that behavior. Those regulated technologies in turn influence or constrain the targeted behavior differently. They "influence the development of code."[23] They are regulations of code that in turn make behavior more regulable.

The question that began this chapter was whether there were similar ways that the government might regulate code on the Internet to make behavior on the Net more regulable. The answer is obviously yes. There are many steps the government might take to make behavior on the network more regulable, and there are obvious reasons for taking those steps.

If done properly, these steps would reduce and isolate untraceable Internet behavior. That in turn would increase the probability that bad behavior would be detected. Increased detection would significantly reduce the expected return from maliciousness. For some significant range of malevolent actors, that shift would drive their bad behavior elsewhere.

This would not work perfectly, of course. No effort of control could ever be perfect in either assuring traceability or tracking misbehavior. But perfection is not the standard. The question is whether the government could put enough incentives into the mix of the network to induce a shift towards traceability as a default. For obvious reasons, again, the answer is yes.

The General Form

If the government's aim is to facilitate traceability, that can be achieved by attaching an identity to actors on the network. One conceivable way to do that would be to require network providers to block actions by individuals not displaying a government-issued ID. That strategy, however, is unlikely, as it is politically impossible. Americans are antsy enough about a national identity card;[24] they are not likely to be interested in an Internet identity card.

But even if the government can't *force* cyber citizens to carry IDs, it is not difficult to create strong *incentives* for individuals to carry IDs. There is no requirement that all citizens have a driver's license, but you would find it very hard to get around without one, even if you do not drive. The government does not require that you keep state-issued identification on your person, but if you want to fly to another city, you must show at least one form of it. The point is obvious: Make the incentive to carry ID so strong that it tips the normal requirements of interacting on the Net.

In the same way, the government could create incentives to enable digital IDs, not by regulating individuals directly but by regulating intermediaries. Intermediaries are fewer, their interests are usually commercial, and they are ordinarily pliant targets of regulation. ISPs will be the "most important and obvious" targets—"focal points of Internet control."[25]

Consider first the means the government has to induce the spread of "digital IDs." I will then describe more what these "digital IDs" would have to be.

First, government means:

- Sites on the Net have the ability to condition access based on whether someone carries the proper credential. The government has the power to require sites to impose this condition. For example, the state could require that gambling sites check the age and residency of anyone trying to use the site. Many sites could be required to check the citizenship of potential users, or any number of other credentials. As more and more sites complied with this requirement, individuals would have a greater and greater incentive to carry the proper credentials. The more credentials they carried, the easier it would be to impose regulations on them.[26]

- The government could give a tax break to anyone who filed his or her income tax with a proper credential.
- The government could impose a 10 percent Internet sales tax and then exempt anyone who purchased goods with a certificate that authenticated their state of residence; the state would then be able to collect whatever local tax applied when it was informed of the purchase.[27]
- The government could charge users for government publications unless they gained access to the site with a properly authenticated certificate.
- As in other Western democracies, the government could mandate voting[28]—and then establish Internet voting; voters would come to the virtual polls with a digital identity that certified them as registered.
- The government could make credit card companies liable for the full cost of any credit card or debit card online fraud whenever the transaction was processed without a qualified ID.
- The government could require the establishment of a secure registry of e-mail servers that would be used to fight spam. That list would encourage others to begin to require some further level of authentication before sending e-mail. That authentication could be supplied by a digital ID.

The effect of each of these strategies would be to increase the prevalence of digital IDs. And at some point, there would be a tipping. There is an obvious benefit to many on the Net to be able to increase confidence about the entity with whom they are dealing. These digital IDs would be a tool to increase that confidence. Thus, even if a site permits itself to be accessed without any certification by the user, any step beyond that initial contact could require carrying the proper ID. The norm would be to travel in cyberspace with an ID; those who refuse would find the cyberspace that they could inhabit radically reduced.

The consequence of this tipping would be to effectively stamp every action on the Internet—at a minimum—with a kind of digital fingerprint. That fingerprint—at a minimum—would enable authorities to trace any action back to the party responsible for it. That tracing—at a minimum—could require judicial oversight before any trace could be effected. And that oversight—at a minimum—could track the ordinary requirements of the Fourth Amendment.

At a minimum. For the critical part in this story is not that the government could induce an ID-rich Internet. Obviously it could. Instead, the important question is the kind of ID-rich Internet the government induces.

Compare two very different sorts of digital IDs, both of which we can understand in terms of the "wallet" metaphor used in Chapter 4 to describe the evolving technology of identity that Microsoft is helping to lead.

One sort of ID would work like this: Every time you need to identify yourself, you turn over your wallet. The party demanding identification rummages through the wallet, gathering whatever data he wants.

The second sort of ID works along the lines of the Identity Layer described in Chapter 4: When you need to identify yourself, you can provide the minimal identification necessary. So if you need to certify that you're an American, only that bit gets revealed. Or if you need to certify that you're over 18, only that fact gets revealed.

On the model of the second form of the digital ID, it becomes possible to imagine then an ultra-minimal ID—an identification that reveals nothing on its face, but facilitates traceability. Again, a kind of digital fingerprint which is meaningless unless decoded, and, once decoded, links back to a responsible agent.

These two architectures stand at opposite ends of a spectrum. They produce radically different consequences for privacy and anonymity. Perfect anonymity is possible with neither; the minimal effect of both is to make behavior traceable. But with the second mode, that traceability itself can be heavily regulated. Thus, there should be no possible traceability when the only action at issue is protected speech. And where a trace is to be permitted, it should only be permitted if authorized by proper judicial action. Thus the system would preserve the capacity to identify who did what when, but it would only realize that capacity under authorized circumstances.

The difference between these two ID-enabled worlds, then, is all the difference in the world. And critically, which world we get depends completely upon the values that guide the development of this architecture. ID-type 1 would be a disaster for privacy as well as security. ID-type 2 could radically increase privacy, as well as security, for all except those whose behavior can legitimately be tracked.

Now, the feasibility of the government effecting either ID depends crucially upon the target of regulation. It depends upon there being an entity responsible for the code that individuals use, and it requires that these entities can be effectively regulated. Is this assumption really true? The government may be able to regulate the telephone companies, but can it regulate a diversity of code writers? In particular, can it regulate code writers who are committed to resisting precisely such regulation?

In a world where the code writers were the sort of people who governed the Internet Engineering Task Force[29] of a few years ago, the answer is probably no. The underpaid heroes who built the Net have ideological reasons to resist government's mandate. They were not likely to yield to its threats. Thus, they would provide an important check on the government's power over the architectures of cyberspace.

But as code writing becomes commercial—as it becomes the product of a smaller number of large companies—the government's ability to regulate it increases. The more money there is at stake, the less inclined businesses (and their backers) are to bear the costs of promoting an ideology.

The best example is the history of encryption. From the very start of the debate over the government's control of encryption, techies have argued that such regulations are silly. Code can always be exported; bits know no borders. So the idea that a law of Congress would control the flow of code was, these people argued, absurd.

The fact is, however, that the regulations had a substantial effect. Not on the techies—who could easily get encryption technologies from any number of places on the Net—but on the businesses writing software that would incorporate such technology. Netscape or IBM was not about to build and sell software in violation of U.S. regulations. The United States has a fairly powerful threat against these two companies. As the techies predicted, regulation did not control the flow of bits. But it did quite substantially inhibit the development of software that would use these bits.[30]

The effect has been profound. Companies that were once bastions of unregulability are now becoming producers of technologies that facilitate regulation. For example, Network Associates, inheritor of the encryption program PGP, was originally a strong opponent of regulation of encryption; now it offers products that facilitate corporate control of encryption and recovery of keys.[31] Key recovery creates a corporate back door, which, in many contexts, is far less restricted than a governmental back door.

Cisco is a second example.[32] In 1998 Cisco announced a router product that would enable an ISP to encrypt Internet traffic at the link level—between gateways, that is.[33] But this router would also have a switch that would disable the encryption of the router data and facilitate the collection of unencrypted Internet traffic. This switch could be flipped at the government's command; in other words, the data would be encrypted only when the government allowed it to be.

The point in both cases is that the government is a player in the market for software. It affects the market both by creating rules and by purchasing products. Either way, it influences the supply of commercial software providers who exist to provide what the market demands.

Veterans of the early days of the Net might ask these suppliers, "How could you?"

"It's just business," is the obvious reply.

EAST COAST AND WEST COAST CODES

Throughout this section, I've been speaking of two sorts of code. One is the "code" that Congress enacts (as in the tax code or "the U.S. Code"). Congress passes an endless array of statutes that say in words how to behave. Some statutes direct people; others direct companies; some direct bureaucrats. The technique is as old as government itself: using commands to control. In our country, it is a primarily East Coast (Washington, D.C.) activity. Call it "East Coast Code."

The other is the code that code writers "enact"—the instructions imbedded in the software and hardware that make cyberspace work. This is code in its modern sense. It regulates in the ways I've begun to describe. The code of Net95, for example, regulated to disable centralized control; code that encrypts regulates to protect privacy. In our country (MIT excepted), this kind of code writing is increasingly a West Coast (Silicon Valley, Redmond) activity. We can call it "West Coast Code."

West Coast and East Coast Code can get along perfectly when they're not paying much attention to each other. Each, that is, can regulate within its own domain. But the story of this chapter is "When East Meets West": what happens when East Coast Code recognizes how West Coast Code affects regulability, and when East Coast Code sees how it might interact with West Coast Code to induce it to regulate differently.

This interaction has changed. The power of East Coast Code over West Coast Code has increased. When software was the product of hackers and individuals located outside of any institution of effective control (for example, the University of Illinois or MIT), East Coast Code could do little to control West Coast Code.[34] But as code has become the product of companies, the power of East Coast Code has increased. When commerce writes code, then code can be controlled, because commercial entities can be controlled. Thus, the power of East over West increases as West Coast Code becomes increasingly commercial.

There is a long history of power moving west. It tells of the clash of ways between the old and the new. The pattern is familiar. The East reaches out to control the West; the West resists. But that resistance is never complete. Values from the East become integrated with the West. The new takes on a bit of the old.

That is precisely what is happening on the Internet. When West Coast Code was born, there was little in its DNA that cared at all about East Coast Code concerns. The Internet's aim was end-to-end communication. Regulation at the middle was simply disabled.

Over time, the concerns of East Coast Coders have become much more salient. Everyone hates the pathologies of the Internet—viruses, ID theft, and spam, to pick the least controversial. That universal hatred has warmed West Coast Coders to finding a remedy. They are now primed for the influence East Coast Code requires: adding complements to the Internet architecture that will bring regulability to the Net.

Now, some will continue to resist my claim that the government can effect a regulable Net. This resistance has a common form: Even if architectures of identification emerge, and even if they become common, there is nothing to show that they will become universal, and nothing to show that at any one time they could not be evaded. Individuals can always work around these technologies of identity. No control that they could effect would ever be perfect.

True. The control of an ID-rich Internet would never be complete. There will always be ways to escape.

But there is an important fallacy lurking in the argument: Just because perfect control is not possible does not mean that effective control is not possible. Locks can be picked, but that does not mean locks are useless. In the context of the Internet, even partial control would have powerful effects.

A fundamental principle of bovinity is operating here and elsewhere. Tiny controls, consistently enforced, are enough to direct very large animals. The controls of a certificate-rich Internet are tiny, I agree. But we are large animals. I think it is as likely that the majority of people would resist these small but efficient regulators of the Net as it is that cows would resist wire fences. This is who we are, and this is why these regulations work.

So imagine the world in which we all could simply establish our credentials simply by looking into a camera or swiping our finger on a thumbprint reader. In a second, without easily forgotten passwords, or easily forged authentication, we get access to the Net, with all of the attributes that are ours, reliably and simply assertable.

What will happen then? When you can choose between remembering a pass-phrase, typing it every time you want access to your computer, and simply using your thumb to authenticate who you are? Or if not your thumb, then your iris, or whatever body part turns out to be cheapest to certify? When it is easiest simply to give identity up, will anyone resist?

If this is selling your soul, then trust that there are truly wonderful benefits to be had. Imagine a world where all your documents exist on the Internet in a "virtual private network," accessible by you from any machine on the Net and perfectly secured by a biometric key.[35] You could sit at any machine, call up your documents, do your work, answer your e-mail, and move on—everything perfectly secure and safe, locked up by a key certified by the markings in your eye.

This is the easiest and most efficient architecture to imagine. And it comes at (what some think) is a very low price—authentication. Just say who you are, plug into an architecture that certifies facts about you, give your identity away, and all this could be yours.

Z-THEORY

"So, like, it didn't happen, Lessig. You said in 1999 that commerce and government would work together to build the perfectly regulable net. As I look through my spam-infested inbox, while my virus checker runs in the background, I wonder what you think now. Whatever was possible hasn't happened. Doesn't that show that you're wrong?"

So writes a friend to me as I began this project to update *Code* v1. And while I never actually said anything about *when* the change I was predicting would happen, there is something in the criticism. The theory of *Code* v1 is missing a part: Whatever incentives there are to push in small ways to the perfectly regulable Net, the theory doesn't explain what would motivate the final push. What gets us over the tipping point?

The answer is not fully written, but its introduction was published this year. In May 2006, the *Harvard Law Review* gave Professor Jonathan Zittrain (hence "Z-theory") 67 pages to explain "The Generative Internet."[36] The article is brilliant; the book will be even better; and the argument is the missing piece in *Code* v1.

Much of *The Generative Internet* will be familiar to readers of this book. General-purpose computers plus an end-to-end network, Zittrain argues, have produced an extraordinarily innovative ("generative") platform for invention. We celebrate the good stuff this platform has produced. But we (I especially) who so celebrate don't pay enough attention to the bad. For the very same design that makes it possible for an Indian immigrant to invent HoTMaiL, or Stanford dropouts to create Google, also makes it possible for malcontents and worse to create viruses and worse. These sorts use the generative Internet to generate evil. And as Zittrain rightly observes, we've just begun to see the evil this malware will produce. Consider just a few of his examples:

- In 2003, in a test designed to measure the sophistication of spammers in finding "open relay" servers through which they could send their spam undetected, within 10 hours spammers had found the server. Within 66 hours they had sent more than 3.3 million messages to 229,468 people.[37]
- In 2004, the Sasser worm was able to compromise more than 500,000 computers—in just 3 days.[38] The year before, the Slammer worm infected 90 percent of a particular Microsoft server—in just 15 minutes.[39]
- In 2003, the SoBig.F e-mail virus accounted for almost 70 percent of the e-mails sent while it was spreading. More than 23.2 million messages were sent to AOL users alone.[40]

These are of course not isolated events. They are instead part of a growing pattern. As the U.S. Computer Emergency Readiness Team calculates, there has been an explosion of security incidents reported to CERT. Here is the graph Zittrain produced from the data:[41]

Number of Security Incidents Reported to CERT/CC, 1988-2003

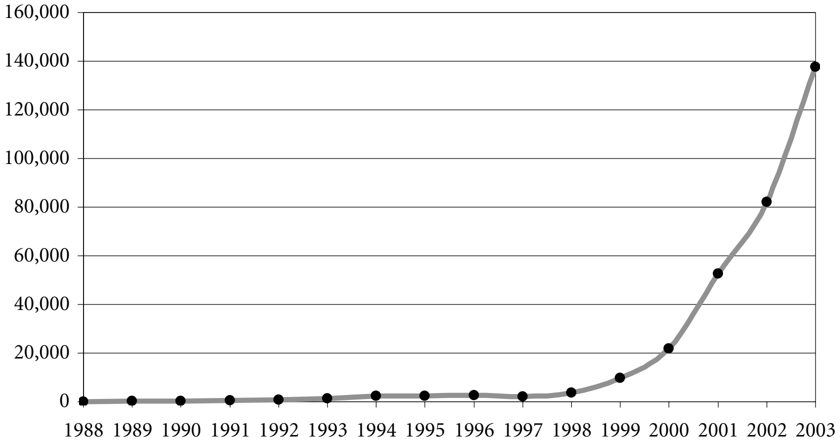

The graph ends in 2004 because CERT concluded that the incidents were so "commonplace and widespread as to be indistinguishable from one another."[42]

That there is malware on the Internet isn't surprising. That it is growing isn't surprising either. What is surprising is that, so far at least, this malware has not been as destructive as it could be. Given the ability of malware authors to get their malicious code on many machines very quickly, why haven't more tried to do real harm?

For example, imagine a worm that worked itself onto a million machines, and in a synchronized attack, simultaneously deleted the hard drive of all million machines. Zittrain's point is not that this is easy, but rather, that it is just as difficult as the kind of worms that are already successfully spreading themselves everywhere. So why doesn't one of the malicious code writers do real damage? What's stopping cyber-Armageddon?

The answer is that there's no good answer. And when there's no good explanation for why something hasn't happened yet, there's good reason to worry that it will happen. And when this happens—when a malware author produces a really devastatingly destructive worm—that will trigger the political resolve to do what so far governments have not done: push to complete the work of transforming the Net into a regulable space.

This is the crucial (and once you see it, obvious) insight of Z-theory. Terror motivates radical change. Think about, for example, the changes in law enforcement (and the protection of civil rights) effected by the "Patriot Act."[43] This massively extensive piece of legislation was enacted 45 days after the terror attacks on 9/11. But most of that bill had been written long before 9/11. The authors knew that until there was a serious terrorist attack, there would be insufficient political will to change law enforcement significantly. But once the trigger of 9/11 was pulled, radical change was possible.

The same will be true of the Internet. The malware we've seen so far has caused great damage. We've suffered this damage as annoyance rather than threat. But when the Internet's equivalent of 9/11 happens—whether sponsored by "terrorists" or not—annoyance will mature into political will. And that political will will produce real change.

Zittrain's aim is to prepare us for that change. His powerful and extensive analysis works through the trade-offs we could make as we change the Internet into something less generative. And while his analysis is worthy of a book of its own, I'll let him write it. My goal in pointing to it here is to provide an outline to an answer that plugs the hole in the theory of *Code* v1. *Code* v1 described the means. Z-theory provides the motive.

There was an awful movie released in 1996 called *Independence Day.* The story is about an invasion by aliens. When the aliens first appear, many earthlings are eager to welcome them. For these idealists, there is no reason to assume hostility, and so a general joy spreads among the hopeful across the globe in reaction to what before had seemed just a dream: really cool alien life.

Soon after the aliens appear, however, and well into the celebration, the mood changes. Quite suddenly, Earth's leaders realize that the intentions of these aliens are not at all friendly. Indeed, they are quite hostile. Within a very short time of this realization, Earth is captured. (Only Jeff Goldblum realizes what's going on beforehand, but he always gets it first.)

My story here is similar (though I hope not as awful). We have been as welcoming and joyous about the Net as the earthlings were about the aliens in *Independence Day;* we have accepted its growth in our lives without questioning its final effect. But at some point, we too will come to see a potential threat. We will see that cyberspace does not guarantee its own freedom but instead carries an extraordinary potential for control. And then we will ask: How should we respond?

I have spent many pages making a point that some may find obvious. But I have found that, for some reason, the people for whom this point should be most important do not get it. Too many take this freedom as nature. Too many believe liberty will take care of itself. Too many miss how different architectures embed different values, and that only by selecting these different architectures—these different codes—can we establish and promote our values.

Now it should be apparent why I began this book with an account of the rediscovery of the role for self-government, or control, that has marked recent history in post-Communist Europe. Market forces encourage architectures of identity to facilitate online commerce. Government needs to do very little—indeed, nothing at all—to induce just this sort of development. The market forces are too powerful; the potential here is too great. If anything is certain, it is that an architecture of identity will develop on the Net—and thereby fundamentally transform its regulability.

But isn't it clear that government should do something to make this architecture consistent with important public values? If commerce is going to define the emerging architectures of cyberspace, isn't the role of government to ensure that those public values that are not in commerce's interest are also built into the architecture?

Architecture is a kind of law: It determines what people can and cannot do. When commercial interests determine the architecture, they create a kind of privatized law. I am not against private enterprise; my strong presumption in most cases is to let the market produce. But isn't it absolutely clear that there must be limits to this presumption? That public values are not exhausted by the sum of what IBM might desire? That what is good for America Online is not necessarily good for America?

Ordinarily, when we describe competing collections of values, and the choices we make among them, we call these choices "political." They are choices about how the world will be ordered and about which values will be given precedence.

Choices among values, choices about regulation, about control, choices about the definition of spaces of freedom—all this is the stuff of politics. Code codifies values, and yet, oddly, most people speak as if code were just a question of engineering. Or as if code is best left to the market. Or best left unaddressed by government.

But these attitudes are mistaken. Politics is that process by which we collectively decide how we should live. That is not to say it is a space where we collectivize—a collective can choose a libertarian form of government. The point is not the substance of the choice. The point about politics is process. Politics is the process by which we reason about how things ought to be.

Two decades ago, in a powerful trilogy drawing together a movement in legal theory, Roberto Unger preached that "it's all politics."[44] He meant that we should not accept that any part of what defines the world is removed from politics—everything should be considered "up for grabs" and subject to reform.

Many believed Unger was arguing that we should put everything up for grabs all the time, that nothing should be certain or fixed, that everything should be in constant flux. But that is not what he meant.

His meaning was instead just this: That we should interrogate the necessities of any particular social order and ask whether they are in fact necessities, and we should demand that those necessities justify the powers that they order. As Bruce Ackerman puts it, we must ask of every exercise of power: Why?[45] Perhaps not exactly at the moment when the power is exercised, but sometime.

"Power," in this account, is just another word for constraints that humans can do something about. Meteors crashing to earth are not "power" within the domain of "it's all politics." Where the meteor hits is not politics, though the consequences may well be. Where it hits, instead, is nothing we can do anything about.

But the architecture of cyberspace is power in this sense; how it is could be different. Politics is about how we decide, how that power is exercised, and by whom.

If code is law, then, as William Mitchell writes, "control of code is power": "For citizens of cyberspace, . . . code . . . is becoming a crucial focus of political contest. Who shall write that software that increasingly structures our daily lives?"[46] As the world is now, code writers are increasingly lawmakers. They determine what the defaults of the Internet will be; whether privacy will be protected; the degree to which anonymity will be allowed; the extent to which access will be guaranteed. They are the ones who set its nature. Their decisions, now made in the interstices of how the Net is coded, define what the Net is.

How the code regulates, who the code writers are, and who controls the code writers—these are questions on which any practice of justice must focus in the age of cyberspace. The answers reveal how cyberspace is regulated. My claim in this part of the book is that cyberspace is regulated by its code, and that the code is changing. Its regulation is its code, and its code is changing.

We are entering an age when the power of regulation will be relocated to a structure whose properties and possibilities are fundamentally different. As I said about Russia at the start of this book, one form of power may be destroyed, but another is taking its place.

Our aim must be to understand this power and to ask whether it is properly exercised. As David Brin asks, "If we admire the Net, should not a burden of proof fall on those who would change the basic assumptions that brought it about in the first place?"[47]

These "basic assumptions" were grounded in liberty and openness. An invisible hand now threatens both. We need to understand how.

One example of the developing struggle over cyber freedoms is the still-not-free China. The Chinese government has taken an increasingly aggressive stand against behavior in cyberspace that violates real-space norms. Purveyors of porn get 10 years in jail. Critics of the government get the same. If this is the people's republic, this is the people's tough love.

To make these prosecutions possible, the Chinese need the help of network providers. And local law requires that network providers in China help. So story after story now reports major network providers—including Yahoo! and Microsoft—helping the government do the sort of stuff that would make our Constitution cringe.

The extremes are bad enough. But the more revealing example of the pattern I'm describing here is Google. Google is (rightly) famous for its fantastic search engine. Its brand has been built on the idea that no irrelevant factor controls its search results. Companies can buy search words, but their results are bracketed and separate from the main search results. The central search results—that part of the screen your eyes instinctively go to—are not to be tampered with.

Unless the company seeking to tamper with the results is China, Inc. For China, Google has promised to build a special routine.[48] Sites China wants to block won't appear in the Google.CN search engine. No notice will be presented. No system will inform searchers that the search results they are reading have been filtered by Chinese censors. Instead, to the Chinese viewer, this will look like normal old Google. And because Google is so great, the Chinese government knows most will be driven to Google, even if Google filters what the government doesn't want its people to have.

Here is the perfect dance of commerce with government. Google can build the technology the Chinese need to make China's regulation more perfectly enabled, and China can extract that talent from Google by mandating it as a condition of being in China's market.

The value of that market is thus worth more to Google than the value of its "neutral search" principle. Or at least, it better be, if this deal makes any sense.

My purpose here is not to criticize Google—or Microsoft, or Yahoo! These companies have stockholders; maximizing corporate value is their charge. Were I running any of these companies, I'm not sure I would have acted differently.

But that in the end is my point: Commerce has a purpose, and government can exploit that to its own end. It will, increasingly and more frequently, and when it does, the character of the Net will change.

Radically so.