

ARCHITECTING INNOVATION

*Lawrence Lessig**

We get trained in ways of seeing things, and then we see in only those ways. We see just what those ways reveal—nothing more, nothing inconsistent. Our world is carved up according to these formulae. Beyond this vision, we see nothing.

We are a long way into this history of our future. It has been a decade since notions of a World Wide Web were first explored, five years since the dominant OS manufacturer built the Web into its OS. The Internet has millions of users—ordinary users—who spend large chunks of their lives in cyberspace. Commerce has found cyberspace and has spent billions to understand how cyberspace might fit into real space economies.

Yet, we are not far into understanding how this future will work. We still look at these questions of power and international influence with the same lens that our fathers used. We ask the same types of questions about sovereignty and international policy. The same actors have the same significance—nothing has changed in the story we tell, save the calendar and the government.

My aim this morning is to sketch a different view. It is to offer a way of understanding how cyberspace has changed things and, hence, a way of speaking about the “Internet and Power.” There is a revolution in Internet relations. It is a revolution in the things being related as power is exercised. And theorists of power and its place in constituting justice need to come to see this revolution and the character of the difference it brings.

Governments think about regulation. They think, in part, about their power to bring about behaviors they like and their power to stop behavior they do not like. They think about this power exercised first through law—through the self-conscious exercise of a threat of force as an inducement to carry into effect a certain set of behaviors.

But, obviously many things regulate behavior—law is just one. We might think about how norms, as well as law, regulate behavior. I spent three glorious years in Cambridge. These were three years learning about how different English norms might be and how severely one might be punished for tripping across different norms. Norms regulate in a way that is similar to law—they punish deviations *ex post*, but norms are nonetheless different from law—their punishment is imposed by a community, not a state.

* Copyright © 2001 by Lawrence Lessig. All rights reserved.

So, too, can we think about how markets regulate. Markets regulate through price. The market conditions access to a resource upon the offering of something of value in exchange—wages for labor, a quid for a quo. The constraint is simultaneous. It is enforced through a community. Deviations from the constraint get remedied through appeals to law and norms.

But, the regulator on which we must focus to understand what is different about cyberspace is what I will call “architecture”—the regulation affected by how the physical world is, whether how it is, how it was made, or how it is found. Architecture regulates behavior. Its constraints are simultaneous, but its constraints get enforced not through the will of the state or through the will of a community but through the physical power of a context or environment.

The idea that architecture might regulate is nothing new; it is certainly not new with cyberspace. David Hackett Fisher describes the founders of New England meticulously laying out the towns they would found so that the relationship of the buildings to each other and to the town square would assure that behavior within the town would be properly regulated. Jeremy Bentham famously described the design of a prison such that all cells would be viewable from one central position so that prisoners would never know whether they were being watched, but that they always could be watched and, so, they would be properly regulated. Napoleon the III had Paris rebuilt so that the boulevards would be broad, making it hard for revolutionaries to blockade the city so that Parisians would be properly regulated. Robert Moses built highway bridges along the roads to the beaches in Long Island so that buses could not pass under the bridges, thereby assuring that only those with cars—mainly white people—would use certain public beaches and that those without cars—largely African Americans—would be driven to use other beaches, so that social relations would be properly regulated.

In all these cases, the constraints of a context are changing or regulating behavior. In each, these constraints are relied upon to regulate this behavior as much as, if differently than, law.

Architecture, law, norms, and markets together regulate behavior. Together they set the terms on which one is free to act or not. Together they set the constraints that affect what is and is not possible. They are four modalities of regulation. They, together, determine how individuals and states within their scope are regulated.

Now, it is important to think about these modalities together because these modalities together compete—one can enable the other; the other might undermine the one. A market might be strengthened by a set of business norms. A set of social norms might be undermined by a market.

Cyberspace presents a particularly virulent interaction. Its architecture, that is, interacts strongly with these other modalities. Depending upon its design, cyberspace can enable the power of social norms, or depending upon its

design, it can disable that power. Depending upon its design, cyberspace can enable a market, or depending upon its design, it can make market functions too costly. And depending upon its design, cyberspace can enable state regulation, or depending upon its design, it can make behavior in cyberspace “unregulable.”

Depending upon its design. For here is the feature on which those of us who want to understand power in this space must focus. Cyberspace is an architecture first. It is a platform that gets designed. It is constituted by a set of code—by software and hardware that makes cyberspace as it is. This code imbeds certain values: it enables certain practices and sets the terms on which life in cyberspace is lived as crucially as the laws of nature set the terms on which life in real space is lived.

Now, most think about this architecture—this code that defines cyberspace—as given. Most think about this code as if it is simply defined, as if it has a nature, and that nature cannot change. As if God gave us cyberspace, and we must simply learn how it is. Most think this and, hence, then spout truths about cyberspace. Slogans that are said to reflect its nature. Cyberspace, it is said, cannot be regulated, behavior there cannot be controlled. Governments are disabled by the nature of that space.

But the first point to see is that this view about cyberspace’s nature is not a view about nature. If anything is socially constructed, cyberspace is. This view about cyberspace’s nature is simply a report about a particular design. A report, that is, about the design that cyberspace had at just the time that cyberspace was coming into the public’s view.

This initial design disabled control. In particular, this initial design disabled two kinds of control. This initial design disabled control by governments, thereby creating a certain kind of liberty in the space. And this initial design disabled control by other actors, or competitors, thereby creating a certain kind of competition and innovation in cyberspace. Certain features of the space made it so. Certain features of its design. Let us focus, first, upon the liberty creating features—the aspects of the design that empowered what we think of as the original freedom of the original Net.

The original Net, for example, protected free speech. Its architecture protected free speech. It protected free speech because one was easily able to say what one wanted without that speech being controlled by others. China could not censor criticisms of China. News of terror in Bosnia could flow freely outside state borders.

The original Net also protected privacy. One could surf on the Net without others knowing who you were. One could live life relatively anonymously. The protocols supported relaying technologies. These technologies made it easy to hide who one was. One could, therefore, be whomever one wanted or no one at all.

The original Net also protected the free flow of content. Perfect copies of text, images, and music could be made for free. These copies could be distributed anywhere on the Net. Content flowed freely.

And, finally, the original Net protected individuals from local regulation. Governments could not condition life on the Net upon conforming to any rule. Governments could not distinguish local from non-local activity. So, behavior here was outside the reach of local government regulation. It was free of control in this way as well.

Now, each of these original liberties of the early Net rested upon a certain feature of the original Net. These freedoms hung upon a certain design. Because it was hard to identify someone, because it was hard to identify where someone came from, because it was hard to identify the content of any Net transaction, because all of this knowledge could not be easily known, behavior in cyberspace could not be regulated.

But what if this feature of the original Net changed? What if it became easier to identify someone? What if it became easier to know where someone came from? What if content became more easily identifiable? What would happen to the original liberty of the Net if these features changed? Well, clearly, if it were easier to identify who someone was or where they came from, and if it were easier to identify content on the Net and from where it originated, then each of these liberties would become vulnerable:

Speech would be less free because dissenters could be tracked.

Privacy would be less secure because identities could be known.

Content would flow less freely because control could be imposed.

Behavior would be more regulated because local states could condition access upon conforming to certain behavior.

That's how things could be—and how things could be is how they are becoming. For the single most important change in the character of cyberspace over the past five years is the erosion of the conditions that made this initial liberty possible. Technologies are being layered onto the original architecture of the Web that change this original design—architectures that make it easier to identify who someone is, architectures that make it easier to know from where they come from, architectures that make it simpler to control the content that they use.

These technologies are essentially private right now. They are tools that commerce is building to make it easier to know who the customer is and to make it easier to control what the customer does with the content she is given. But the same technologies of identification and tracking would also make it easier for governments to regulate. The same technologies would make it easier for governments to control.

You have heard about these technologies in many different contexts. At their most crude, they are the cookie technologies that make it easy for Web sites to track you. In a more sophisticated case, they are digital certificates that will collect personal information about you and make it possible for you to certify this information to others. And just around the corner, there are biometric technologies that will make it possible to certify who you are without any real risk. These technologies will make identification possible. And this identification will, in turn, make regulation possible. And this regulation, in turn, will change the character of the original Net from liberty to something else.

Consider one example. iCraveTV was an Internet broadcaster in Canada. Under Canadian law, they were permitted to capture the broadcasts from Canadian television and rebroadcast that in any medium they wanted. iCraveTV decided to rebroadcast that TV across the Internet. Now, free TV is not allowed in the United States. Under United States law, the rebroadcaster must negotiate with the original broadcaster. So, iCraveTV used technologies to block Americans from getting access to iCraveTV. Canadians were to get access to free TV, Americans were not. But it is the nature of the existing architecture of the Net that it is hard to control perfectly who gets access to what. So, there were a number of Americans who were able to get access to iCraveTV, despite the company's efforts to block foreigners.

Hollywood did not like this much. So, as quickly as you could say "cut," it had filed a lawsuit in a Pittsburgh federal court, asking that court to shut down the Canadian site. The argument was this: whether or not free TV is legal in Canada, it is not legal in the United States. And, so, since some in the United States might, God forbid, get access to free TV, the United States court should shut down free TV. Copyright laws in the United States were being violated. Massive and quick response by the federal courts was called for.

Now, step back for a moment and think about the equivalent claim being made elsewhere. Imagine, for example, a German court entering a judgment against Amazon.com, ordering Amazon.com to stop selling *Mein Kampf* "anywhere" because someone in Germany had succeeded in accessing *Mein Kampf* from Amazon.com. Or imagine a court in China ordered an American ISP to shut down its dissidents' site because the speech at issue was illegal in China. It would take just a second for an American to say that these suits violate the concept of free speech on the Net; that they undermine the free flow of information; that they are an improper extension of state power into the world of cyberspace.

But free speech did not register in this Pittsburgh court. The idea of the rights of Canadians to their free TV did not matter. The court ordered the site shut down until the site could prove that it could keep non-Canadians out. iCraveTV promised to try. And it quickly developed technologies that would succeed in zoning cyberspace based on geography.

Now, the pattern here should be clear. Though nations like the United States will sing about the importance of free speech in cyberspace and about keeping cyberspace free, when it comes to issues of national security—as all things copyright are—values fall away. The push will be to zone the space, to allow rules to be imposed that are local. And the technologies for zoning will quickly develop.

The future will be different from this past. It will be a future that enables control that the past has not seen; it will be a future of regulation, imposed locally, by local governments. These regulations will be carried into effect by rules, built into technologies, built into the architecture of the space. The character of the original space—its liberty—will change.

But it is a second kind of control that I want to focus on here, not the control by government over behavior in cyberspace. That story is becoming fairly well known. Instead, I want to focus on the control the space allows by actors within the space on other actors within the space. And to see the point, I want to begin with a story told by John Naughton in his extraordinary book *A Brief History of the Future*.

In 1964, a Rand Researcher named Paul Baran proposed to the Defense Department a design for a telecommunication network that was very much like the design of the current Internet. It was not quite the architecture of the Internet, and Baran was probably not the first to propose such a design. But the idea was radical and important enough that the Defense Department asked their network experts to comment on the design. Their experts were AT&T. AT&T did not like the plan. As AT&T executive Jack Osterman said of the plan, “First it can’t possibly work, and if it did, damned if we are going to allow the creation of a competitor to ourselves.”

Allow.

The telephone network had a particular architecture. That architecture embedded certain principles. Those principles were that the network owner—AT&T—got to decide how the network would be used. The network centralized that decision, and this centralized design was supported by the regulations of the FCC. Until the late 1960s, and partially continuing until the breakup of AT&T in 1984, the network owner had the power to decide what kinds of innovations would be *allowed* on the telecommunications network. The architecture embedded this power to decide.

This principle affected innovation. Innovators knew that before their ideas about how a telecommunications network should be used and would be adopted, AT&T would have to approve their ideas. They knew their ideas would need the *permission* of someone else before they would run, and they knew that this someone else had an interest in the existing model of telecommunications. Some new ideas would be consistent with that model—no doubt they would be embraced. But other new ideas would be inconsistent with this model. They had

a snowball's chance in hell. Any rational innovator—or at least, those with a bottom line to support—would turn their innovative energies elsewhere.

At the core of the original design of the Internet is a different architectural principle. This principle has a different effect on innovation. First described by network architects Jerome Saltzer, David P. Reed, and David Clark in 1981, this principle, called the “end-to-end” argument, guides network designers in placing intelligence in the network at the ends and to keep the network, itself, stupid. Stupid networks, smart applications.

While this principle was first described in terms of efficiency, it soon became clear that it entailed an important corollary. This is the principle of competitive neutrality. What end-to-end meant was that the network was not in a position to discriminate. It was not capable of deciding which kinds of applications should run or what forms of content should be permitted. The network was stupid—it processed packets blindly. It could no more decide what packets were “competitors” than the post office could determine which letters criticize it.

This architecture has an effect on innovation. It encouraged innovation. Innovators knew that if they designed a new application or new form of content, the network would run it. They did not have to negotiate with every possible network before they could run their programs on the network. Even if the new application challenged the dominant network application, the network would run it. The test of success, therefore, was not whether the innovation fit with the business model of the network owner, the test of success was whether the market demanded it. End-to-end, thus, architected innovation. It is a principle, a value, that was built into the original Net and had the consequence of guaranteeing extraordinary innovation and creativity on the original Net.

Now end-to-end is not the only principle of the Net's original architecture. It is not the only way in which innovation is encouraged. Other structures dominant in the original design do very much the same thing. Much of the original Net was open source. Open source software is software that leaves its source code available for others to take and use and develop as they wish. That means two things: that the resources are available for others to use towards innovation, and that the system—because it is open and not in the control of any one company—cannot discriminate against new uses. Like the principle of end-to-end, the platform of an open source system remains neutral. It is a neutral platform that invites innovation elsewhere.

A neutral platform. A type of commons that all can draw upon. A common resource that produces a common good. A heritage of neutrality that has produced the greatest innovation and creativity that the world has seen.

This neutrality, this commons, gets built by a certain design of technology that governs the Net and a certain way that the code of the Net lives. This neutrality is a principle, chosen.

This principle of neutrality is changing. We are changing it. As the network moves from narrowband to broadband, it is being architected in a way that violates end-to-end. It is being built in a way that gives the network owners power over the uses of the network that will be allowed. It is being constructed to give the owners of conduit—cable and wireless access—the power to say what kinds of content flow on the Net. It is being built to empower strategic action by network owners, built to undermine the neutrality of the original regime.

In the United States, this is happening in the context of broadband cable. The principle players are AT&T, which is buying up all the cable monopolies it can, and now AOL, which has purchased Time-Warner. Both companies were, at one time, in favor of an open and neutral platform on the Net. But once each became the owner of cable lines themselves, each was interested in being able to control what went across those cable lines.

For example, video. Cable monopolies stream video to television sets. The Internet (remember iCraveTV) enables the streaming of video to computers. But cable in the United States has restricted the streaming of video to computers—while charging for the streaming of video to TVs. And when asked whether video would ever be permitted to computers on the cable lines, the head of AT&T's Internet services said, "We didn't spend 56 billion on a cable network to have the blood sucked from our veins."

The point is the principle—the principle that we are giving up. By allowing these network owners to architect the network in violation of end-to-end, we are allowing them to architect away the incentives for innovation that the original Net created. The choice about the code is a choice about the innovation that code encourages.

Now, my point in raising these examples is to direct your attention to what is important here. In both liberty and innovation, cyberspace is changing. In both contexts, an initial freedom and opportunity is being given up. Principles from the original Net are being eroded. The future will not be as the past.

These changes are occurring because the architecture of the Net is changing. And we are changing that architecture. The architecture is developing to reduce the liberty and reduce the innovation. It is developing to make cyberspace like real space—regulated and concentrated, controlled and bland.

In both cases, the source of this change is commerce. The source is an influence that resists the unregulated and seeks protected monopoly power. This source is having its effect.

But as we think about power and its effect, as we think about the state using power to have a different effect, as we think about policy, internationally and locally, we must have a different focus. Rather than focusing exclusively on the shifts of power among nations or the shifts of power between some nations in particular, our focus should first be upon the shift in power that the rise of this

architecture called cyberspace creates. How it, now, represents a source of power, and how the character of that power gets determined by its design. And then second, we should focus on how the particular character of its original power is now changing. How we are allowing it to change as commerce alters its code.

Struggles of power get played out upon this stage set by the architecture of the space—whether states have power, and how much; whether competitors have power over other competitors, and how much. An account of international relations that ignored this stage would be as incomplete an account as one that ignored China.

But when we come to see its place, its role in regulating these relations, its potential for regulating differently, we will then see how its original regulation is changing. How we are allowing it to change. How we stand back and watch as this ecology of innovation gets remade, undermined, killed, by changes to the architecture that defined a free space, a commons, that has built the greatest revolution in creativity we have seen.

