Foreign Policy
1779 Massachusetts Ave., NW,
Washington, DC 20036-2103
Phone: 202-939-2230
Fax: 202-483-4430

## The Internet Under Siege

**Who owns the Internet? Until recently, nobody. That's because, although the Internet was "Made in the U.S.A.," its unique design transformed it into a resource for innovation that anyone in the world could use. Today, however, courts and corporations are attempting to wall off portions of cyberspace. In so doing, they are destroying the Internet's potential to foster democracy and economic growth worldwide.**

By Lawrence Lessig

The Internet revolution has ended just as surprisingly as it began. None expected the explosion of creativity that the network produced; few expected that explosion to collapse as quickly and profoundly as it has. The phenomenon has the feel of a shooting star, flaring unannounced across the night sky, then disappearing just as unexpectedly. Under the guise of protecting private property, a series of new laws and regulations are dismantling the very architecture that made the Internet a framework for global innovation.

Neither the appearance nor disappearance of this revolution is difficult to understand. The difficulty is in accepting the lessons of the Internet's evolution. The Internet was born in the United States, but its success grew out of notions that seem far from the modern American ideals of property and the market. Americans are captivated by the idea, as explained by Yale Law School professor Carol Rose, that the world is best managed "when divided among private owners" and when the market perfectly regulates those divided resources. But the Internet took off precisely because core resources were not "divided among private owners." Instead, the core resources of the Internet were left in a "commons." It was this commons that engendered the extraordinary innovation that the Internet has seen. It is the enclosure of this commons that will bring about the Internet's demise.

This commons was built into the very architecture of the original network. Its design secured a right of decentralized innovation. It was this "innovation commons" that produced the diversity of creativity that the network has seen within the United States and, even more dramatically, abroad. Many of the Internet innovations we now take for granted (not the least of which is the World Wide Web) were the creations of "outsiders"—foreign inventors who freely roamed the commons. Policymakers need to understand the importance of this architectural design to the innovation and creativity of the original network. The potential of the Internet has just begun to be realized, especially in the developing world, where many "real space"

alternatives for commerce and innovation are neither free nor open.



Yet old ways of thinking are reasserting themselves within the United States to modify this design. Changes to the Internet's original core will in turn threaten the network's potential everywhere—staunching the opportunity for innovation and creativity. Thus, at the moment this transformation could have a meaningful effect, a counterrevolution is succeeding in undermining the potential of this network.

The motivation for this counterrevolution is as old as revolutions themselves. As Niccolò Machiavelli described long before the Internet, "Innovation makes enemies of all those who prospered under the old regime, and only lukewarm support is forthcoming from those who would prosper under the new." And so it is today with us. Those who prospered under the old regime are threatened by the Internet. Those who would prosper under the new regime have not risen to defend it against the old; whether they will is still a question. So far, it appears they will not.

**The Neutral Zone**

A "commons" is a resource to which everyone within a relevant community has equal access. It is a resource that is not, in an important sense, "controlled." Private or state-owned property is a controlled resource; only as the owner specifies may that property be used. But a commons is not subject to this sort of control. Neutral or equal restrictions may apply to it (an entrance fee to a park, for example) but not the restrictions of an owner. A commons, in this sense, leaves its resources "free."

Commons are features of all cultures. They have been especially important to cultures outside the United States—from communal tenure systems in Switzerland and Japan to irrigation communities within the Philippines. But within American intellectual culture, commons are treated as imperfect resources. They are the object of "tragedy," as ecologist Garrett Hardin famously described. Wherever a commons exists, the aim is to enclose it. In the American psyche, commons are unnecessary vestiges from times past and best removed, if possible.

For most resources, for most of the time, the bias against commons makes good sense. When resources are left in common, individuals may be driven to overconsume, and therefore deplete, them. But for some resources, the bias against commons is blinding. Some resources are not subject to the "tragedy of the commons" because some resources cannot be "depleted." (No matter how much we use Einstein's theories of relativity or copy Robert Frost's poem "New Hampshire," those resources will survive.) For these resources, the challenge is to induce provision, not to avoid depletion. The problems of provision are very different from the problems of depletion—confusing the two only leads to misguided policies.

This confusion is particularly acute when considering the Internet. At the core of the Internet is a design (chosen without a clear sense of its consequences) that was new among large-scale computer and communications networks. Named the "end-to-end argument" by network theorists Jerome Saltzer, David Clark, and David Reed in 1984, this design influences where "intelligence" in the network is placed. Traditional computer-communications systems located intelligence, and hence control, within the network itself. Networks were "smart"; they were designed by people who believed they knew exactly what the network would be used for.

But the Internet was born at a time when a different philosophy was taking shape within computer science. This philosophy ranked humility above omniscience and anticipated that network designers would have no clear idea about all the ways the network could be used. It therefore counseled a design that built little into the network itself, leaving the network free to develop as the ends (the applications) wanted.

The motivation for this new design was flexibility. The consequence was innovation. Because innovators needed no permission from the network owner before different applications or content got served across the network, innovators were freer to develop new modes of connection. Technically, the network achieved this design simply by focusing on the delivery of packets of data, oblivious to either the contents of the packets or their owners. Nor does the network concern itself that all the packets make their way to the other

side. The network is "best efforts"; anything more is provided by the applications at both ends. Like an efficient post office (imagine!), the system simply forwards the data along.

Since the network was not optimized for any single application or service, the Internet remained open to new innovation. The World Wide Web is perhaps the best example. The Web was the creation of computer scientist Tim Berners-Lee at the European Organization for Nuclear Research (CERN) laboratory in Geneva in late 1990. Berners-Lee wanted to enable users on a network to have easy access to documents located elsewhere on the network. He therefore developed a set of protocols to enable hypertext links among documents located across the network. Because of end-to-end, these protocols could be layered on top of the initial protocols of the Internet. This meant the Internet could grow to embrace the Web. Had the network compromised its commitment to end-to-end—had its design been optimized to favor telephony, for example, as many in the 1980s wanted—then the Web would not have been possible.

This end-to-end design is the "core" of the Internet. If we can think of the network as built in layers, then the end-to-end design was created by a set of protocols implemented at the middle layer—what we might call the logical, or code layer, of the Internet. Below the code layer is a physical layer (computers and the wires that link them). Above the code layer is a content layer (material that gets served across the network). Not all these layers were organized as commons. The computers at the physical layer are private property, not "free" in the sense of a commons. Much of the content served across the network is protected by copyright. It, too, is not "free."

At the code layer, however, the Internet is a commons. By design, no one controls the resources for innovation that get served across this layer. Individuals control the physical layer, deciding whether a machine or network gets connected to the Internet. But once connected, at least under the Internet's original design, the innovation resources for the network remained free.

No other large scale network left the code layer free in this way. For most of the history of telephone monopolies worldwide, permission to innovate on the telephone platform was vigorously controlled. In the United States in 1956, AT&T successfully persuaded the U.S. Federal Communications Commission to block the use of a plastic cup on a telephone receiver, designed to block noise from the telephone microphone, on the theory that AT&T alone had the right to innovation on the telephone network.

The Internet might have remained an obscure tool of government-backed researchers if the telephone company had maintained this control. The Internet would never have taken off if ordinary individuals had been unable to connect to the network by way of Internet service providers (ISPs) through already existing telephone lines. Yet this right to connect was not preordained. It is here that an accident in regulatory history played an important role. Just at the moment the Internet was emerging, the telephone monopoly was being moved to a different regulatory paradigm. Previously, the telephone monopoly was essentially free to control its wires as it wished. Beginning in the late 1960s, and then more vigorously throughout the 1980s, the government began to require that the telephone industry behave neutrally—first by insisting that telephone companies permit customer premises equipment (such as modems) to be connected to the network, and then by requiring that telephone companies allow others to have access to their wires.

This kind of regulation was rare among telecommunications monopolies worldwide. In Europe and throughout the world, telecommunications monopolies were permitted to control the uses of their networks. No requirement of access operated to enable competition. Thus no system of competition grew up around these other monopolies. But when the United States broke up AT&T in 1984, the resulting companies no longer had the freedom to discriminate against other uses of their lines. And when ISPs sought access to the local Bell lines to enable customers to connect to the Internet, the local Bells were required to grant access equally. This enabled a vigorous competition in Internet access, and this competition meant that the network could not behave strategically against this new technology. In effect, through a competitive market, an end-to-end design was created at the physical layer of the telephone network, which meant that an end-to-end design could be layered on top of that.

This innovation commons was thus layered onto a physical infrastructure that, through regulation, had important commons-like features. Common-carrier regulation of the telephone system assured that the system could not discriminate against an emerging competitor, the Internet. And the Internet itself was created, through its end-to-end design, to assure that no particular application or use could discriminate against any other innovations. Neutrality existed at the physical and code layer of the Internet.

An important neutrality also existed at the content layer of the Internet. This layer includes all the content streamed across the network—Web pages, MP3s, e-mail, streaming video—as well as application programs that run on, or feed, the network. These programs are distinct from the protocols at the code layer, collectively referred to as TCP/IP (including the protocols of the World Wide Web). TCP/IP is dedicated to the public domain.

But the code above these protocols is not in the public domain. It is, instead, of two sorts: proprietary and nonproprietary. The proprietary includes the familiar Microsoft operating systems and Web servers, as well as programs from other software companies. The nonproprietary includes open source and free software,

especially the Linux (or GNU/Linux) operating system, the Apache server, as well as a host of other plumbing-oriented code that makes the Net run.

Nonproprietary code creates a commons at the content layer. The commons here is not just the resource that a particular program might provide—for example, the functionality of an operating system or Web server. The commons also includes the source code of software that can be drawn upon and modified by others. Open source and free software ("open code" for short) must be distributed with the source code. The source code must be free for others to take and modify. This commons at the content layer means that others can take and build upon open source and free software. It also means that open code can't be captured and tilted against any particular competitor. Open code can always be modified by subsequent adopters. It, therefore, is licensed to remain neutral among subsequent uses. There is no "owner" of an open code project.

In this way, and again, parallel to the end-to-end principle at the code layer, open code decentralizes innovation. It keeps a platform neutral. This neutrality in turn inspires innovators to build for that platform because they need not fear the platform will turn against them. Open code builds a commons for innovation at the content layer. Like the commons at the code layer, open code preserves the opportunity for innovation and protects innovation against the strategic behavior of competitors. Free resources induce innovation.

**An Engine of Innovation**

The original Internet, as it was extended to society generally, mixed controlled and free resources at each layer of the network. At the core code layer, the network was free. The end-to-end design assured that no network owner could exercise control over the network. At the physical layer, the resources were essentially controlled, but even here, important aspects were free. One had the right to connect a machine to the network or not, but telephone companies didn't have the right to discriminate against this particular use of their network. And finally, at the content layer, many of the resources served across the Internet were controlled. But a crucial range of software building essential services on the Internet remained free. Whether through an open source or free software license, these resources could not be controlled.

This balance of control and freedom produced an unprecedented explosion in innovation. The power, and hence the right, to innovate was essentially decentralized. The Internet might have been an American invention, but creators from around the world could build upon this network platform. Significantly, some of the most important innovations for the Internet came from these "outsiders."

As noted, the most important technology for accessing and browsing the Internet (the World Wide Web) was not invented by companies specializing in network access. It wasn't America Online (AOL) or Compuserve. The Web was developed by a researcher in a Swiss laboratory who first saw its potential and then fought to bring it to fruition. Likewise, it wasn't existing e-mail providers who came up with the idea of Web-based e-mail. That was co-created by an immigrant to the United States from India, Sabeer Bhatia, and it gave birth to one of the fastest growing communities in history—Hotmail.

And it wasn't traditional network providers or telephone companies that invented the applications that enabled online chatting to take off. The original community-based chatting service (ICQ) was the invention of an Israeli, far from the trenches of network design. His service could explode (and then be purchased by AOL for $400 million) only because the network was left open for this type of innovation.

Similarly, the revolution in bookselling initiated by Amazon.com (through the use of technologies that "match preferences" of customers) was invented far from the traditional organs of publishers. By gathering a broad range of data about purchases by customers, Amazon—drawing upon technology first developed at MIT and the University of Minnesota to filter Usenet news—can predict what a customer is likely to want. These recommendations drive sales, but without the high cost of advertising or promotion. Consequently, booksellers such as Amazon can outcompete traditional marketers of books, which may account for the rapid expansion of Amazon into Asia and Europe.

These innovations are at the level of Internet services. Far more profound have been innovations at the level of content. The Internet has not only inspired invention, it has also inspired publication in a way that would never have been produced by the world of existing publishers. The creation of online archives of lyrics and chord sequences and of collaborative databases collecting information about compact discs and movies demonstrates the kind of creativity that was possible because the right to create was not controlled.

Again, the innovations have not been limited to the United States. OpenDemocracy.org, for example, is a London-based, Web-centered forum for debate and exchange about democracy and governance throughout the world. Such a forum is possible only because no coordination among international actors is needed. And it thrives because it can engender debate at a low cost.

This history should be a lesson. Every significant innovation on the Internet has emerged outside of traditional providers. The new grows away from the old. This trend teaches the value of leaving the

platform open for innovation. Unfortunately, that platform is now under siege. Every technological disruption creates winners and losers. The losers have an interest in avoiding that disruption if they can. This was the lesson Machiavelli taught, and it is the experience with every important technological change over time. It is also what we are now seeing with the Internet. The innovation commons of the Internet threatens important and powerful pre-Internet interests. During the past five years, those interests have mobilized to launch a counterrevolution that is now having a global impact.

This movement is fueled by pressure at both the physical and content layers of the network. These changes, in turn, put pressure on the freedom of the code layer. These changes will have an effect on the opportunity for growth and innovation that the Internet presents. Policymakers keen to protect that growth should be skeptical of changes that will threaten it. Broad-based innovation may threaten the profits of some existing interests, but the social gains from this unpredictable growth will far outstrip the private losses, especially in nations just beginning to connect.

**Fencing Off the Commons**

The Internet took off on telephone lines. Narrowband service across acoustic modems enabled millions of computers to connect through thousands of ISPs. Local telephone service providers had to provide ISPs with access to local wires; they were not permitted to discriminate against Internet service. Thus the physical platform on which the Internet was born was regulated to remain neutral. This regulation had an important effect. A nascent industry could be born on the telephone wires, regardless of the desires of telephone companies.

But as the Internet moves from narrowband to broadband, the regulatory environment is changing. The dominant broadband technology in the United States is currently cable. Cable lives under a different regulatory regime. Cable providers in general have no obligation to grant access to their facilities. And cable has asserted the right to discriminate in the Internet service it provides.

Consequently, cable has begun to push for a different set of principles at the code layer of the network. Cable companies have deployed technologies to enable them to engage in a form of discrimination in the service they provide. Cisco, for example, developed "policy-based routers" that enable cable companies to choose which content flows quickly and which flows slowly. With these, and other technologies, cable companies will be in a position to exercise power over the content and applications that operate on their networks.

This control has already begun in the United States. ISPs running cable services have exercised their power to ban certain kinds of applications (specifically, those that enable peer-to-peer service). They have blocked particular content (advertising from competitors, for example) when that content was not consistent with their business model. The model for these providers is the model of cable television generally—controlling access and content to the cable providers' end.

The environment of innovation on the original network will change according to the extent that cable becomes the primary mode of access to the Internet. Rather than a network that vests intelligence in the ends, the cable-dominated network will vest an increasing degree of intelligence within the network itself. And to the extent it does this, the network will increase the opportunity for strategic behavior in favor of some technologies and against others. An essential feature of neutrality at the code layer will have been compromised, reducing the opportunity for innovation worldwide.

Far more dramatic, however, has been the pressure from the content layer on the code layer. This pressure has come in two forms. First, and most directly related to the content described above, there has been an explosion of patent regulation in the context of software. Second, copyright holders have exercised increasing control over new technologies for distribution.

The changes in patent regulation are more difficult to explain, though the consequence is not hard to track. Two decades ago, the U.S. Patent Office began granting patents for software-like inventions. In the late 1990s, the court overseeing these patents finally approved the practice and approved their extension to "business methods." The European Union (EU), meanwhile, initially adopted a more skeptical attitude toward software patents. But pressure from the United States will eventually bring the EU into alignment with American policy.

In principle, these patents are designed to spur innovation. But with sequential and complementary innovation, little evidence exists that suggests such patents will do any good, and there is increasing evidence that they will do harm. Like any regulation, patents tax the innovative process generally. As with any tax, some firms—large rather than small, U.S. rather than foreign—are better able to bear that tax than others. Open code projects, in particular, are threatened by this trend, as they are least able to negotiate appropriate patent licenses.

The most dramatic restrictions on innovation, however, have come at the hands of copyright holders. Copyright is designed to ensure that artists control their "writings" for a limited time. The aim is to secure to

copyright holders a sufficient interest to produce new work. But copyright laws were crafted in an era long before the Internet. And their effect on the Internet has been to transfer control over innovation in distribution from many innovators to a concentrated few.

The clearest example of this effect is online music. Before the Internet, the production and distribution of music had become extraordinarily concentrated. In 2000, for example, five companies controlled 84 percent of music distribution in the world. The reasons for this concentration are many—including the high costs of promotion—but the effect of concentration on artist development is profound. Very few artists make any money from their work, and the few that do are able to do so because of mass marketing from record labels. The Internet had the potential to change this reality. Both because the costs of distribution were so low, and because the network also had the potential to significantly lower the costs of promotion, the cost of music could fall, and revenues to artists could rise.

Five years ago, this market took off. A large number of online music providers began competing for new ways to distribute music. Some distributed MP3s for money (eMusic.com). Some built technology for giving owners of music easier access to their music (mp3.com). And some made it much easier for ordinary users to "share" their music with other users (Napster). But as quickly as these companies took off, lawyers representing old media succeeded in shutting them down. These lawyers argued that copyright law gave the holders (some say hoarders) of these copyrights the exclusive right to control how they get used. American courts agreed.

To keep this dispute in context, we should think about the last example of a technological change that facilitated a much different model for distributing content: cable TV, which has been accurately hailed as the first great Napster. Owners of cable television systems essentially set up antenna and "stole" over-the-air broadcasts and then sold that "stolen property" to their customers. But when U.S. courts were asked to stop this "theft," they refused. Twice the U.S. Supreme Court held that this use of someone else's copyrighted material was not inconsistent with copyright law.

When the U.S. Congress finally got around to changing the law, it struck an importantly illustrative balance. Congress granted copyright owners the right to compensation from the use of their material on cable broadcasts, but cable companies were given the right to broadcast the copyrighted material. The reason for this balance is not hard to see. Copyright owners certainly are entitled to compensation for their work. But the right to compensation shouldn't translate into the power to control innovation. Rather than giving copyright holders the right to veto a particular new use of their work (in this case, because it would compete with over-the-air broadcasting), Congress assured copyright owners would get paid without having the power to control—compensation without control.

The same deal could have been struck by Congress in the context of online music. But this time, the courts did not hesitate to extend control to the copyright holders. So the concentrated holders of these copyrights were able to stop the deployment of competing distributors. And Congress was not motivated to respond by granting an equivalent compulsory right. The aim of the recording company's strategy was plain enough: shut down these new and competing models of distribution and replace them with a model for distributing music online more consistent with the traditional model.

This trend has been supported by the actions of Congress. In 1998, Congress passed the Digital Millennium Copyright Act (DMCA), which (in)famously banned technologies designed to circumvent copyright protection technologies and also created strong incentives for ISPs to remove from their sites any material claimed to be a violation of copyright.

On the surface both changes seem sensible enough. Copyright protection technologies are analogous to locks. What right does anyone have to pick a lock? And ISPs are in the best position to assure that copyright violations don't occur on their Web sites. Why not create incentives for them to remove infringing copyrighted material?

But intuitions here mislead. A copyright protection technology is just code that controls access to copyrighted material. But that code can restrict access more effectively (and certainly less subtly) than copyright law does. Often the desire to crack protection systems is nothing more than a desire to exercise what is sometimes called a fair-use right over the copyrighted material. Yet the DMCA bans that technology, regardless of its ultimate effect.

More troubling, however, is that the DMCA effectively bans this technology on a worldwide basis. Russian programmer Dimitry Sklyarov, for example, wrote code to crack Adobe's eBook technology in order to enable users to move eBooks from one machine to another and to give blind consumers the ability to "read" out loud the books they purchased. The code Sklyarov wrote was legal where it was written, but when it was sold by his company in the United States, it became illegal. When he came to the United States in July 2001 to talk about that code, the FBI arrested him. Today Sklyarov faces a sentence of 25 years for writing code that could be used for fair-use purposes, as well as to violate copyright laws.

Similar trouble has arisen with the provision that gives ISPs the incentive to take down infringing

copyrighted material. When an ISP is notified that material on its site violates copyright, it can avoid liability if it removes the material. As it doesn't have any incentive to expose itself to liability, the ordinary result of such notification is for the ISP to remove the material. Increasingly, companies trying to protect themselves from criticism have used this provision to silence critics. In August 2001, for example, a British pharmaceutical company invoked the DMCA in order to force an ISP to shut down an animal rights site that criticized the British company. Said the ISP, "It's very clear [the British company] just wants to shut them up," but ISPs have no incentive to resist the claims.

In all these cases, there is a common pattern. In the push to give copyright owners control over their content, copyright holders also receive the ability to protect themselves against innovations that might threaten existing business models. The law becomes a tool to assure that new innovations don't displace old ones—when instead, the aim of copyright and patent law should be, as the U.S. Constitution requires, to "promote the progress of science and useful arts."

These regulations will not only affect Americans. The expanding jurisdiction that American courts claim, combined with the push by the World Intellectual Property Organization to enact similar legislation elsewhere, means that the impact of this sort of control will be felt worldwide. There is no "local" when it comes to corruption of the Internet's basic principles. As these changes weaken the open source and free software movements, countries with the most to gain from a free and open platform lose. Those affected will include nations in the developing world and nations that do not want to cede control to a single private corporation. And as content becomes more controlled, nations that could otherwise benefit from vigorous competition in the delivery and production of content will also lose. An explosion of innovation to deliver MP3s would directly translate into innovation to deliver telephone calls and video content. Lowering the cost of this medium would dramatically benefit nations that still suffer from weak technical infrastructures.

Policymakers around the world must recognize that the interests most strongly protected by the Internet counterrevolution are not their own. They should be skeptical of legal mechanisms that enable those most threatened by the innovation commons to resist it. The Internet promised the world—particularly the weakest in the world—the fastest and most dramatic change to existing barriers to growth. That promise depends on the network remaining open to innovation. That openness depends upon policy that better understands the Internet's past.

*Lawrence Lessig is professor of law at Stanford University. He is author of* The Future of Ideas: The Fate of the Commons in the Connected World *(New York: Random House, 2001) and* Code and Other Laws of Cyberspace *(New York: Basic Books, 1999).*

Want to Know More?