Vanderbilt Journal of Entertainment & Technology Law

Volume 1 | Issue 1 Article 4

1999

The Architecture of Privacy: Remaking Privacy in Cyberspace

Lawrence Lessig

Remaking Privacy in Cyberspace

By Lawrence Lessig

There is a part of anyone's life that is *monitored*, and there is a part that can be *searched*. The monitored is that part of one's day-to-day life that others see, that others notice, that others could take note of, and respond to, if response were appropriate. The searchable is the part of one's life that leaves, or is, a record. As I walk down the street, my behavior is monitored. If I walked down the street in a small village in Mainland China, my behavior would be monitored quite extensively. This monitoring in both cases would be transitory—people would notice if I were walking with an elephant, or walking in a dress; but if there were nothing special about my walk, if I simply blended into the crowd, then I might be noticed for the moment, but I would be forgotten soon after. The searchable is less transitory. The scribblings in my diary leave a record of my thoughts. They can be searched. Things in my house are a record of what I possess. It too can be searched.

The Architecture of

And the recordings on my telephone answering machine are a record of who called, and what they said. It can be searched as well. These parts of my life don't pass so easily away. They are not in the same way ephemeral. They instead remain to be reviewed if there is interest, and if technology, and the law, would permit.

This is an essay about privacy. My aim is to understand privacy through these two very different ideas. Privacy, in the sense that I mean here, can be described by these two different ideas. It stands in competition with these ideas. It is that part of life that is left after one subtracts, as it were, the monitored and the searchable. A life where less is monitored is a life where more is private; and life where less can (legally or technologically) be searched is also a life where more is private. By understanding the technologies of these two different ideas, the monitored and the searchable—understanding, as it were, their *architectures*—we understand something of the privacy that any particular social context makes possible.

These contexts are many. They differ dramatically across the world. But in this essay, I want to use this notion of the monitored and the searchable to compare privacy across contexts, and to see just why the context we are about to enter is so extraordinarily different from any we have known.

For my claim is that we are entering an age when privacy will be fundamentally altered—an age when the extent of the monitored, and the reach of searchable, is far greater than anything we have known thus far. We can choose to let this change occur, or we can choose to do something in response. After making plain the kind of change we should expect, my aim is to make understandable a range of responses, and to argue, if only implicitly, in favor of one particular response within that range.

THE MONITORED AND THE SEARCHABLE

The monitored, as I described it, is that part of one's life that is watched. It is the part that is watched in an ordinary or regular way. My focus here is not the infrequent spy, though if spying became extensive enough, spying would be part of the monitored. Nor is it the periodic patrol of police. The monitored, as I mean it, is the regular and persistent watching of people or machines, whether the behavior watched is considered "public" or not.

Monitoring in social life is quite familiar. It is life in the small town. People living in a relatively small community, known by their neighbors, monitored as they come and go, as they buy in the market, as they associate at a local pub. Everything in that life, it is said, is seen. Everything in that life, it is said, was therefore known by others. In that world, it is said, one could not build the modern liberal conception of privacy. Privacy was what went on in one's head, not in one's life.

This is the picture that Americans often have of America at the founding. And it is the picture that leads many to say that there was no concept of "privacy" in America at the founding. Life then was life in public. One lived in small towns, everyone knew one's neighbors, and everyone knew one's business. If you stayed out too late, or if you drank too much, or if you associated with the wrong people, or if you were rude to another in public—if you in any way breached an elaborate set of norms about how citizens were to behave, your breach would be noticed, and you would suffer the consequences of the breach. The social norms of such a society regulated individuals in that society, and they could therefore regulate much of the individual's life in such a society—since much of an individual's life was, in this sense, public, or in my terms, monitored.

But this type of monitoring—the monitor of the small town, or the monitor of the community—has important features that we should not overlook. The first is its relative transience; the second is who is doing the monitoring. My neighbors might remember that I was at the local market Saturday morning; they may even remember with whom I was talking; but they are not likely to remember exactly at what time I was there, or everyone with whom I spoke. Nor will they know what I bought, or how much I paid, or whether I paid with large denomination bills or small. Of course, and again, if I did something out of the ordinary—if I brought my elephant to the market, or came with a woman who wasn't my wife—then my actions in a small town might be noticed in a less transitory way. Then my actions might be remembered. But in the ordinary case, they are not remembered. They are monitored for the moment, and then the record from that monitoring is forgotten. It is erased.

Privacy

More important than its transience, however, is the character of who is doing the monitoring. Small towns, of course, have their busybodies—people who pry into the business of others—and they have their moral prudes—people whose standards are much stricter than most. But these enforcers of community norms are on the fringe. They define the extreme of a much narrower core. And it this core of moral ideals that sets the limits on freedom that a community might define. The monitoring of a community serves this core; but its limit sets the limit on the burden of this monitoring. To an outsider, these norms might seem harsh. They might seem wrong. But for members of that community, they are just the sort of norms that the "ordinary" in that community obey. For them, the norms are not extreme or selective. They are not easily manipulated or changed. They are a set of influences that apply generally to like cases. And they get their force because they are applied by the collective, acting as a normative community.

We are entering an age when privacy in any sense of that term will be fundamentally altered—an age when the extent of the monitored, and the reach of the searchable, is far greater than anything we have known thus far.

These are the features of one particular architecture of monitoring. In a moment we will consider other architectures with different features. But before we consider these, consider the other part to privacy's balance—the searchable. And consider it again, if you will, in the context of a small town, or, say, in early America.

As I've defined the term, the searchable is a function not only of what records there are that could be searched, but also of the technologies of searching, and the legal protections against the use of such technologies.

Consider the technologies first. In early America, these technologies were crude. There was no simple way to monitor—to hear, for example, a conversation going on between two people, locked securely in their own house. One might eavesdrop, but not easily and not with great success. And there was no cheap way to search. The searchable—letters, diaries, stuff in my house—was searchable only if the police got access to my property; the law protected me from their wrongful access, and the very nature of the architecture of property protected against wrongful access. The common law, and the architecture of property, combined to establish a zone of privacy that neither the state, nor individuals, could easily breach.

The American Constitution guaranteed this protection of the common law. The Fourth Amendment required that searches be conducted only if reasonable, and that the warrant to search be granted only if there was probable cause to search. This constitutional affirmation of the value of privacy combined with the legal protections of the common law—protections again against trespass, or other invasions of privacy—gave legal support to the technological or architectural support for privacy that existed at the time.

The searchable then is determined by two different factors. The first is the architecture of the social world at the time—at the framing, crude technologies for searching, relatively inefficient means of collecting data. These inefficiencies themselves constituted a kind of a protection; they made it hard to search. And they were supplemented by the protections of law. The law protected individuals against search; it limited the reasons the police could use for searching; it was a second line of defense against the invasion of prying eyes.

Privacy in this original context was then the product of this balance. On the one side, there was a life that was monitored by structures that support social norms. On the other, there was the protection of law, and architecture, that combined to raise the costs of searching. My life on the street might be monitored by my neighbors, but that monitoring produced few searchable records; and those records that were searchable were protected by both the architecture of property—that my walls were not made of glass, or that my door could be double-locked—and by law, both constitutional and as developed by the common law courts. The balance of privacy then was this balance between the monitored and the protections against search.

PRESERVATION ACROSS CONTEXTS

As my story so far should make clear, much about this balance of privacy—at that time, and in any time—

depends upon existing technology. If what softens the burdens of monitoring is that monitoring is relatively transient, then technologies that eliminate transience increase the burden of monitoring. If what constitutes much of the protection of privacy in the home is that one who would breach it must physically enter the home, then technologies that allow invasion without physical invasion are technologies that reduce this privacy. Technologies in both cases can change; the question for law in both cases is how to respond to these changes so that privacy is preserved.

The question was best raised in the Supreme Court in 1928, in the case of Olmstead v. United States.[1] In the midst of America's last great war on drugs—prohibition—the federal government deployed wire-tapping as a device for collecting evidence. State laws forbade wiretapping, and the contracts that telephone companies had with their customers also promised that the wires would not be tapped. Nonetheless, the federal government ignored these protections and invaded the privacy of the defendants' phones. In the case of Olmstead, the defendants challenged that wiretap on the grounds that it violated the Fourth Amendment.

The Supreme Court was not receptive. In its view, that the Fourth Amendment protected against trespass only. Since wiretapping did not involve a trespass, the Fourth Amendment did not protect against it. Hence evidence collected through wiretapping would be admissible to convict Olmstead for violating the laws against prohibition.

Justice Brandeis, however, had a different view—a different view of the Constitution, and a different view about the scope of the Fourth Amendment. Certainly, Brandeis wrote, the Constitution when originally authored protected only against trespass. But when it was authored, trespass was the only effective way to violate someone's privacy. But in 1928 that was no longer the case. In 1928, much of life had already moved onto the wires. And much of private life was now conducted on the telephone. In such a world, Brandeis argued, the protections of the Fourth Amendment should be read to protect privacy on the phone as much as privacy in the home. To protect the same degree of privacy as the framers did, Brandeis argued, it was necessary to protect against more than trespass.

Brandeis' technique should be ours as well. His approach was to first identify values from the original Fourth Amendment, and *translate* those values into the context of cyberspace. Brandeis read beyond the specific applications that the framers had in mind to find the meaning that they intended to constitutionalize. His aim was to carry that meaning of the framers into the context of 1928.

We need the same technique today. We can't help but consider the technologies, or as I've called them, the architectures of privacy in evaluating the world of privacy we are entering. For the world we are entering is about to change these architectures of privacy more completely and more extensively than any similar change in the past. And we can see this change by considering two stories—the first about the monitored; the second, about the searchable.

Peter Lewis, writing in the NEW YORK TIMES, in an article titled "Forget Big Brother," begins his story with the following account:

> Surveillance cameras followed the attractive young blond woman through the lobby of the midtown Manhattan hotel, kept a glassy eye on her as she rode the elevator up to the 23rd floor and peered discreetly down the hall as she knocked at the door to my room. I have not seen the videotapes, but I can imagine the digital readout superimposed on the scenes, noting the exact time of the encounter. That would come in handy if someone were to question later why this woman, who is not my wife, was visiting my hotel room during a recent business trip. The cameras later saw us heading off to dinner and to the theater—a middle aged married man from Texas with his arm around a pretty East Village woman young enough to be his daughter.[2]

"As a matter of fact," Lewis writes, "she is my daughter."

Lewis' is a story of the monitored—a hint of the emerging world of monitoring that is already constituting life in real space, and which promises even greater sway in cyberspace. Add to the cameras the credit card receipts, the telephone logs, the airplane tickets, the toll booths, the check-in records at the hotel, the records from room service—add in all the records that get collected in the ordinary course in life and the scope of monitoring begins to be clear.

Cyberspace will be even worse—or better, depending upon your perspective. Jerry Kang summarizes the difference well:

> [I]magine the following two visits to a mall, one in real space, the other in cyberspace. In

real space, you drive to a mall, walk up and down its corridors, peer into numerous shops, and stroll through corridors of inviting stores. Along the way, you buy an ice-cream cone with cash. You walk into a bookstore and flip through a few magazines. Finally, you stop at a clothing store and buy a friend a silk scarf with a credit card. In this narrative, numerous persons interact with you and collect information along the way. For instance, while walking through the mall, fellow visitors visually collect information about you, if for no other reason than to avoid bumping into you. But such information is general—e.g., it does not pinpoint the geographical location and time of the sighting—is not in a format that can be processed by a computer, is not indexed to your name or another unique identifier, and is impermanent, residing in short-term human memory. You remain a barely noticed stranger. One important exception exists: The scarf purchase generates data that are detailed, computer-processable, indexed by name, and potentially permanent.

By contrast, in cyberspace, the exception becomes the norm: Every interaction is like the credit card purchase. ... In this alternate universe, you are invisibly stamped with a bar code as soon as you venture outside your home. There are entities called "road" providers, who provide the streets and ground you walk on, who track precisely where, when, and how fast you traverse the lands, in order to charge you for your wear on the infrastructure. As soon as you enter the cyber-mall's domain, the mall begins to track you through invisible scanners focused on your bar code. It automatically records which stores you visit, which windows you browse, in which order, and for how long. The specific stores collect even more detailed data when you enter their domain. For example, the cyber-bookstore notes which magazines you skimmed, recording which pages you have seen and for how long, and notes the pattern, if any, of your browsing. It notes that you picked up briefly a health magazine featuring an article on St. John's Wort, read for seven minutes a news weekly detailing a politician's sex scandal, and flipped ever-so-quickly through a tabloid claiming that Elvis lives. Of course, whenever any item is actually purchased, the store as well as the credit, debit, or virtual cash company that provides payment through cyberspace, takes careful notes of what you bought—in this case, a silk scarf, red, expensive.[3]

In both stories, the monitored increases. In both, the scope of one's life subject to monitoring changes. In both cases, this change is made possible by a change in the architecture of each space. The architecture is designed to capture data about our ongoing exchanges and transactions in ordinary life.

The data that these systems collect are much like the data that a community in small town might collect. But again there are important differences. Unlike the data that the community might collect, the information from this monitoring is permanent and searchable. It is not information collected and then discarded (since forgotten); it is information that is collected, kept, and searchable—and not by not just the community but by anyone who wants access to its facts. Think of the billions of gigabytes of e-mail messages stored across the world; or the tapes of telephone records archived by telephone companies; or the archives of frequent flyer miles, or credit card receipts, or calling card debits, or cash machines withdrawals, or toll booth records—think about all these, and you begin to get a sense of the extraordinary data that is coming to be collected as matter of routine, as a matter of what is ordinarily monitored.

This increase in the monitored increases the searchable in two different ways. (1) More can be searched, as more data are collected; (2) searching is cheaper, as the monitored becomes more easy to scan. And this in turn leads, perhaps paradoxically, to an even greater reduction in legal protections against such searches. Consider each change in turn.

Consider each change in turn. The first change is the more familiar, but we should separate its costs into two parts. One part represents the costs borne by the searcher; the other the costs borne by the person being searched. The costs borne by the searcher are those costs involved in executing the search—the time and expenses, etc. The costs borne by the person being searched are not just the subjective costs, but also the intrusion and disruption of the search.

Modern technologies are quickly reducing costs of both kinds. In real space, technologies such as telephoto lenses, long distance microphones, infrared cameras, and body scans, all make it cheaper to detect whatever the searcher is seeking. And in cyberspace, the change is all the more dramatic as data move onto a common protocol network, and systems for data matching become all the more sophisticated. In both cases, the changes will mean a sharp reduction in the costs of a particular search, and hence an increase, in this aspect at least, of the searchable.

Limits, in other words, on searching—both practical and legal—are being eroded. And the result of this erosion will be an ever-increasing range of one's life that it is, at any time in the future, the subject of discovery.

The same change is occurring with the costs borne by the person being searched. For these same devices—devices to scan bodies from a distance, devices to listen through walls from hundreds of feet away, searches of online data which the owner never notices, wiretapping—have become efficient techniques for the searcher and less burdensome for the person being searched.

But it is this second reduction that yields the paradox that I adverted to earlier. For by increasing the efficiency of a search, the changing technologies reduce the legal justifications for interfering with the searches. As searches become more efficient, the scope of "reasonable" searches increases. In the ordinary case, the legal grounds for limiting the power of the state to search have been the burdens imposed on the person being searched. So that as these burdens are removed, there is less and less justification for limiting the state's right to search. Thus as the costs of searching fall, the legal grounds for restricting the search fall as well.

An example will make the point. Searches, the Constitution requires, must be "reasonable." So consider the following. Imagine a worm—a bit of computer code designed to work its way across the net and locate holes in the architecture of the net such that it can place itself onto the hard disks of computer users. The worm is designed not to do any damage. It does not attach itself to any system or application file. The worm instead simply places itself onto a hard disk and searches that disk.

Say this worm were designed by the Federal Bureau of Investigation. And say the worm were designed to search for a particular file—an illegal file, let's say, either a file with a national security document, or an illegal copy of some software code. The worm was designed to search disks without the user noticing; it did its work completely in background. If it found what it was looking for, it would report back to the FBI the location of the file; if it didn't, it would simply destroy itself. The worm would not be able to search beyond its mandate.

Would such a worm violate the constitutional right of privacy? I believe this is a very hard question. Certainly in a sense one might call this an invasion of property, but no longer is the Fourth Amendment tied to conceptions of property. The test under the Fourth Amendment now is simply whether the search is reasonable. Here, the search imposes no burden on the innocent, and only burdens the guilty. It is, in this sense, an efficient search. It is a general search, but because it imposes none of the costs of a general search, it might well be understood best as a reasonable search—like the sniff of a dog at the airport, except here there is not even the fear of the dog.

The worm is just an example, but it points to a more general point. More is being monitored; more can be searched cheaply; more can be searched without imposing any burden on the person being searched—searched efficiently, that is. Limits, in other words, on searching—both practical and legal—are being eroded. And the result of this erosion will be an ever-increasing range of one's life that it is, at any time in the future, the subject of discovery.

How should we understand this change? How should we understand its source? Its source is the change we will see in the architecture of a networked world. In real space, the default is that data are not collected. In real space, it takes effort—either the effort of a community, or the effort of a spy—to gather data. That is the architecture of the real world. And for most of our history, this architecture meant that any data so gathered were, in essence, useless. It was costly to hold, costly to use, and costly to collect.

But the architecture of cyberspace is different. Or

rather, the architecture is quickly becoming different. The architecture of cyberspace can be such that collecting data is the default. The world there can be made such that in the ordinary case, information is collected ceaselessly—invisibly, behind the scenes, efficiently, with no burden on the user. The information is collected; it is more easily searched; and the legal protections against its search—protections grounded in the burden that a search would create—disappear.

This monitoring would be done by the state—by a small group separate from the community. And this separateness is extraordinarily significant, in two very different ways.

And so should we ask: Just how should we respond? How should we respond to this change in technology—to these changes in the architecture of cyberspace that yield a world unlike any we have known before.

The answer is not obvious, but if we put it in a regulatory context, some possibilities might become clear. That is my aim in the next section—to sketch a way of understanding this regulatory context, a model for understanding this problem of regulation. And in the final section, I'll use that model to help explain the differences in the responses of Europe and the United States, and to say something about the possibilities within each.

RESPONDING TO CHANGE

We should keep this issue in perspective. It is not as if the last two hundred years before the Internet were years without technological change. It is not as if we have never faced such questions before. Obviously, the question of individual privacy has been a dominant theme in legal thought for much of modern legal history. And plenty of nations have responded to the changes by enacting legal proscriptions designed to replicate or create protections of an earlier period.

Some nations have, but not the United States. For while most modern democracies have enacted significant legal protections for privacy, we have not. We have been slower to respond and have been much more laissez-faire in our response. We have no general federal statute protecting privacy, whether informational privacy or data privacy. We don't even have federal statutes effectively protecting medical privacy—the only group with that sort of protection is individuals in drug rehab clinics. Instead, where we have responded with laws, the laws are limited to particular problems or contexts. We have very effective protections for data about what videos people rent, but only because a particular prominent American was embarrassed by the publication of the records of the videos he rented. American law is sporadic and partial—incomplete, from the perspective of data privacy in Europe, and inconsequential for most real protections.

The reasons for this lack of law protecting privacy are complex—they relate in part to a general skepticism about legal protection generally; they relate in the balance to the extraordinary lobbying power of interests that would use the data affected by informational privacy regulation. And we should not expect this feature of American law to change dramatically in the short term. Privacy here is not about to be protected by law in the way that privacy in Europe, and parts of Asia, is.

But does that mean that our privacy will remain vulnerable? Or put another way, is law the only kind of protection we might expect? My sense is "no."

Think about the ways in which privacy is protected in real space—the many ways, and not just the protections of law. The law is one protection for privacy, but it is not the only protection, or the most important. Norms protect privacy as well. At least among individuals, norms limit the kinds of questions one might ask, or the kinds of gossip one might listen to. And among corporations, norms restrict the kind of uses that these companies will make of the data they collect. These constraints are different from law—they are enforced, for example, not by the state, but by the sanctions of other members of a particular community. But they are nonetheless a source of constraint, functioning to protect privacy.

The market is another type of protection. Reputation in the market is affected by the use corporations make of privacy data, and in some cases, firms can offer more expensive services with a greater promise of privacy protection.

But in the story I've told so far, the most significant constraint protecting, or possibly eroding, privacy is the constraint of architecture. High walls make secure hous-

es; sophisticated locks keep all but the most skilled burglar out; thick walls can't be listened through; thick curtains don't reveal. All these are features of the architecture of a particular space. And all these features in obvious ways increase, or extend, the privacy of a particular space.

It is against this background, then, that we should consider the state of data privacy today. For I've said already that laws in America are relatively slight and are unlikely to be strengthened anytime soon. But given these alternatives, our question should be whether these alternatives might supplement the law to create a context in which privacy is protected. Do they provide alternatives to the law that might fill the gap that our laissez-faire regime permits?

One alternative, for example, would be norms. This is the solution of the Clinton Administration to the problem of data privacy. The administration wants industry to develop codes for regulating the handing of personal data. It wants industry to develop these codes on its own, and then enforce them without the involvement of the state. Industry would develop its own form of self-regulation, and the state would rely on this self-regulation to protect the privacy of its citizens.[4]

There is much to be skeptical about with this solution —not the least of which being that the interests of commerce might well be different from the interests of the consumer. But it represents an alternative, the effectiveness of which must be considered when accounting for the interests protecting privacy.

A second alternative is architecture—technologies for re-creating privacy where other technologies may have erased it. The most common example here is encryption —especially public key encryption, which would facilitate individuals hiding more effectively facts about themselves that they don't want third parties to know.

But encryption won't hide transactional data—it won't hide the monitoring of clickstreams, or telephone log records. And it won't easily hide records kept about us by third parties—except to the extent those records are protected by others. Moreover, encryption may actually increase the technologies of monitoring and searching, for it facilitates an architecture within which identity can be established, and hence architectures which will require that identity be established. Public key encryption makes it easy to hide what one says. But it also makes it easy to authenticate who one is. Encryption facilitates both hiding and authenticating, for the same technology that locks a conversation can be used to verify an identity. A digital signature, for example, can certify that I sent this, or a digital certificate can certify that I am who I say I am. And it is this second part of the technology for encryption—this part that makes authentication possible—that we should consider when weighing its effect on privacy.

As the cost of authenticating falls, we should expect the use of authenticating technologies to increase. As it is easier to say who I am, we should expect the growth of technologies that ask who I am. The two will work together, for knowing who I am is valuable data. Thus it again will increase the data knowable, in a sense, by the system; it again is an architecture that will advance the ends of monitoring.

For this reason, I don't believe one can say—absolutely, or without qualification—that the development of encryption technologies will increase individual privacy. In the terms that began this essay, encryption may well reduce the searchable, by protecting what I hide; but by reducing the cost of authentication, it might well increase the monitored, and hence increase the searchable again. The technology, like much in this field, is Janus-faced—freedom-enhancing from one perspective, control-enhancing from another.

A better solution, I suggest, is one that links the protection of architecture with the incentives of the market. Information is an asset. It is a resource which has become extremely valuable. And as it has become extremely valuable, commerce has tried to exploit it. This use has a cost—an externality borne by those who would prefer that this data not be used. So the trick is to construct a regime where those who would use the data internalize this cost, by paying those whose data are used.

The laws of property are one such regime. If the law gave individuals the rights to control their data, or more precisely, if those who wanted to use that data had first to secure the right to use it, then a negotiation would occur over whether, and how much, data should be used. The market could negotiate these rights, if a market in these rights could be constructed.

The benefits of a market would be many. Most important among these benefits would be the ability of the market to recognize diversity. A property regime gives the holder of the property right the power to hold out—until the buyer is willing to pay what the seller demands. But what this means is that people can hold out to different degrees. The problem with this property regime, however, is its costs of negotiating the price to

be paid. It would be impossible to imagine dickering with each click on the web. So how could this property regime be created?

It is here that the change in the architecture I alluded to before comes into play. For there are a number of designs that code writers are proposing that might make this structure of negotiation possible.

One example is the regime of the Platform for Privacy Preferences (P3P), designed by the World Wide Web Consortium (W3C). P3P is a standard for negotiating protocols on the web—a standard, that is, for negotiating protocols about privacy. It facilitates setting the terms on which users will enter a site, for example, and then only entering sites that satisfy those terms. In the language of P3P's authors:

> Sites with practices that fall within the range of a user's preference could, at the option of the user, be accessed "seamlessly." Otherwise, users will be notified of a site's practices and have the opportunity to agree to those terms or other terms and continue browsing if they wish.[5]

The web has already made possible person-to-machine communication and person-to-person communication. Architectures like P3P make possible machine-to-machine communication. This means machines can bear the cost of this negotiation and could act as our agents to protect our privacy.

This solution again mixes both a market and architectural response. It is a solution that imagines the two working together to create a kind of protection for privacy that law alone couldn't provide. If successful, it might protect some individual data—not all, and certainly not for all purposes. But some, or perhaps enough, and certainly more than we now have.

LOOKING AHEAD

We are fast entering an age where more can be known, and more efficiently collected, then at any time in our history ever. These changes are brought about by a change in architectures. Of the constraints that might protect privacy, this constraint—architecture—has shifted most significantly. Its shift has an ambiguous quality—it makes possible an efficiency we have not before seen; and it makes likely an extent of monitoring we have not yet known.

One response to this change is law—the response of the Europeans. Laws could be enacted to reconstruct the privacy lost. But there are other responses beyond law—the response of norms, the market, and architecture. I have sketched one that relies on the joint product of two (architecture and the market), and no doubt there are others. The loss of privacy is not inevitable. Responses are possible.

That is the hopeful account. But I want now to end on a note of skepticism, or better, anxiety, about where we are. For as much as we might envision a time when changes could restore a degree of privacy, we should not ignore the changes that are already occurring, and the vulnerability that these changes will create just now.

For the lack of laws protecting data notwithstanding, governments are moving to take advantage of the efficiencies these new architectures facilitate. In Taiwan, for example, the government is developing smart card technologies, combining national insurance information and identity information—including fingerprints—on a single card. These cards will also contain a digital signature, identifying the holder when used with a governmental data base. They are envisioned to be complete records for each individual—perfect identifications, and perfect links with that person's past. Efficient IDs—far better than the IDs we have today.

These efficiencies, of course, are valuable. But they beg for structures that check their use. They beg for structures built into the system that might help assure that they don't become tools of misuse. As a balance to these advances, we must create structures that assure control consistent with values of privacy within our tradition.

A kind of inefficiency should be built into these emerging technologies—an inefficiency that makes it harder for these technologies to be misused. Certainly, it is hard to argue for features of the architecture in cyberspace that will make it more difficult for government to do its work. It is hard to argue that less is more.

But though hard, this is not an argument unknown in the history of constitutional democracies. Indeed, it is the core of much of the design of many of the most successful—that we build into such constitutions structures of restraint that check and limit the efficiency of government, protecting against the tyranny of the majority.

This view helps explain much about the common constitutional rights in a constitutional democracy. They are, as John Perry Barlow has called them, "bugs" in the code of government: elements designed to make government function less efficiently, so that rights are better protected. These "bugs" have value in contexts beyond the context of constitutional rights. They also have value

in the very structure of government itself. One doesn't want a perfectly efficient prosecutor, for fear that the prosecution will grow tyrannical. One doesn't want an unimpeded executive, for fear that the executive will become arbitrary. One doesn't want a perfectly powerful and efficient legislature. One builds into a constitutional democracy limits on effectiveness of governmental power, to protect against abuse of that power.

The architectures of control that are emerging in this cyberworld are not the architectures of control of the traditional community. *Communities* are not, or would not be, monitoring behavior and enforcing norms through self-enforcement. This monitoring would be done by the state—by a small group separate from the community. And this separateness is extraordinarily significant, in two very different ways.

The first difference is size. The "community," however one understands that term, is not the group that is controlling life in this emerging architecture of control. The group that gets the benefit of these architectures of control is the government. Governments, like guns, need not be bad; but when, like guns, they are placed in the wrong hands, they can become quite dangerous. And this is what power through knowledge means: that a small group has a great power, and that therefore the risk of tyranny by this group is all the more great. The rules or requirements that can be enforced by this government are not necessarily the rules or requirements that would be enforced by the community. Its leaders get their power by pretending to enforce the will of the community, but instead enforce whatever will the small group might represent. They can stifle dissent—not because the community necessarily would, but because the architecture of control that has emerged gives them the power to monitor.

The second difference is even more important. If we have learned anything about how communities function —if we have learned anything about the kinds of behavior that support or sustain a community, and the kinds of interventions that destroy it—then we have learned that for a community to sustain itself, the community itself must enforce its rules. The norms of a community are sustained only so long as members of the community are involved in the enforcement of those norms. Norms can not be imposed externally, and in this context governments are often external. If this enforcement is performed by someone else—by the state, or by some other separate enforcing entity—then the community loses the practice of such enforcement and weakens its bonds. Only an inefficient community can sustain itself as a community; an efficient community (one that has institutions to efficiently enforce its norms) would self-destruct. If members don't bear the cost of enforcing the rules of their community, the community will fade.

Privacy needs protection when architectures make more transparent; it gets protection when law and code give individuals greater control. I've described one solution to the problem that changes in the net are now creating. No doubt there are others. But my point is not the particulars. My point is more general. What is missing in discourse about cyberspace and its regulation is a richer understanding of the range of architectures that are possible. We must develop an attitude that analyzes architecture as critically as it analyzes laws—an attitude that understands the politics in both. We will only resolve finally and properly how this world should be made when we understand that we, in this critical sense, are responsible for its making. ◆

This essay is drawn from a lecture delivered at the Taiwan Net '98 conference in Taipei during March 1998.

---

1 277 U.S. 438 (1928).

2 Peter H. Lewis, *Forget Big Brother*, N.Y. TIMES, Mar. 19, 1998, at E1.

3 Jerry Kang, *Information Privacy in Cyberspace Transactions*, 50 STAN. L. REV. 1193, 1198-99 (1998).

4 See H.R. 2368, 105th Cong. (1997).

5 (Oct. 22, 1998) <www.w3c.org/P3P/P3FAQ.html#What is P3.>.