THE DARK SIDE OF INTERNET FREEDOM

THE NET DELUSION

EVGENY MOROZOV

“Evgeny Morozov offers a rare note of wisdom and common sense, on an issue overwhelmed by digital utopians.” —MALCOLM GLADWELL

The NET DELUSION

The Dark Side of Internet Freedom

EVGENY MOROZOV

PUBLICAFFAIRS
New York

Published in the United States by PublicAffairs™, a member of the Perseus Books Group.

Printed in the United States of America.

Designed by Timm Bryson
Text set in 11.75 point Arno Pro

Library of Congress Cataloging-in-Publication Data

Morozov, Evgeny.
The net delusion : the dark side of internet freedom / Evgeny Morozov. — 1st ed.
p. cm.
Includes bibliographical references and index.
ISBN 978-1-58648-874-1 (alk. paper)
1. Internet—Political aspects. 2. Internet—Censorship. 3. Computers—Access control. 4. Freedom of information. I. Title.
HM851.M665 2011
303.48'33—dc22
2010039066

First Edition
10 9 8 7 6 5 4 3 2 1

chapter six

Why the KGB Wants You to Join Facebook

Imagine that you are a target of some deeply mysterious spying operation. While you happily poke your online friends, tweet your breakfast plans, and shop for Christmas presents, all your online activity is being secretly reported to an unknown party. Imagine that someone has also broken into your computer and is using it to launch DDoS attacks. They could be targeting Saudi websites about philosophy or dissident Georgian bloggers. You have no idea that your computer is part of this mysterious cyber-army, let alone who is being attacked or why. It's as if a stranger has been secretly reading your diary and also using it to clobber a passerby.

This is precisely what happened to a number of brave activists from Vietnam who in 2009 were protesting the building of a new bauxite mine in their country. (The project is a joint venture between Chalco, a subsidiary of China's state-run aluminum company Chinalco, and the Vietnamese government.) Their computers were compromised, allowing an unknown third party not only to monitor their online activity but also to attack other online targets in Vietnam and elsewhere. But theirs was not a case of basic computer illiteracy, where pressing the wrong button or visiting a weird porn site could surrender months of hard work to a nasty virus. It's quite likely that the Vietnamese dissidents did no such thing, avoiding any suspicious-looking sites and attachments. What could have gone wrong?

Vietnam, nominally still ruled by a Communist Party, boasts a burgeoning Internet culture, with antigovernment bloggers mounting frequent campaigns about social issues, especially the poorly regulated sprawling urban development. The government, concerned that its tight hold on public life is beginning to loosen, has been trying to reassert control, preferably without drawing much ire from Vietnam's trading partners in the West. The authorities, all too keen on harvesting the information benefits of globalization, do not shy away from computers or the Internet outright. In April 2010 they embarked on an ambitious crusade to supply farmers in more than one thousand communes with free computers, so that they could, as one official put it, "contact and consult with . . . scientists . . . [about] the current epidemics on their breeds and seeds." The government was even kind enough to organize computer training courses for the farmers.

Those opposing the government's paradigm of "modernization at all costs" are unlikely to be invited to attend such courses. In 2009 two of the most vocal blogs that challenged the government, *Bauxite Vietnam* and *Blogosin,* became targets of powerful DDoS attacks similar to those launched against Tomaar and Cyxymu. Soon, *Bauxite Vietnam* was forced down the "digital refugee" route, eventually emerging on a Google-owned blogging service, while the blogger behind *Blogosin* told his readers that he was quitting blogging altogether to "focus on personal matters." Those attacks made it quite clear that the Vietnamese government was up-to-date on the rapidly evolving nature of Internet control and wouldn't stop at just blocking access to particular websites.

Most likely, the antimine activists, careful as they were, inadvertently hit a government trap that allowed the secret police to establish remote control over their computers. And what a trap it was: Someone broke into the server that hosted the website of the Vietnamese Professionals Society (VPS), a trusted diaspora organization, and replaced one of the most popular downloads, a simple computer program that facilitated typing in the Vietnamese language, to an almost identical file—"almost" because it also contained a virus. Anyone who downloaded and installed the software risked turning their computer into a powerful spy and attack hub. Such breaches of security are generally hard to detect, for everything seems to be working normally, and no suspicious activity is taking place.

The Vietnamese activists might have never actually discovered that their every online move was being followed were it not for the buzz generated by the high-profile cyber-attacks that hit Google in December 2009. While investigating the mysterious origins of those sly attacks, researchers from McAffee, a computer security firm, accidentally unearthed the clandestine spying operation in Vietnam and initially believed the two to be related (they were not). The media buzz generated by McAffee's unexpected discovery—also heavily publicized by Google through their own channels—probably generated enough coverage in the Western press to protect the Vietnamese activists from immediate persecution, even if a lot of their private data might have been compromised nevertheless. It's impossible to say how many similar spying operations go undetected, putting authoritarian governments ahead of their opponents.

Never Trust Anyone with a Website

But many such surveillance campaigns—especially when heavily publicized in the media—have effects that extend far beyond the mere gathering of information. Knowing that they might be watched by government agents but not knowing how exactly such surveillance happens, many activists might lean toward self-censorship or even stop engaging in risky online behavior altogether. Thus, even if authoritarian governments cannot actually accomplish what the activists fear, the pervasive climate of uncertainty, anxiety, and fear only further entrenches their power.

Such schemes have much in common with the design of the perfect prison, the panopticon, described by the nineteenth-century British utilitarian philosopher Jeremy Bentham. The point of such systems is to exert control over prisoners' behavior, even when nobody is watching them, by never letting the prisoners know if they are being watched. Governments, of course, are quite happy to overstate their actual capabilities, for such boasting works to their advantage. Thus, in January 2010, when Ahmadi Moghaddam, Iran's police chief, boasted that "the new technologies allow us to identify conspirators and those who are violating the law, without having to control all people individually," he must have known that his words would have an effect even if he had greatly exaggerated his capability. When no one quite knows just how extensive government surveillance really is, every new arrest of a blogger—whether it's based on genuine surveillance practices, tips-off from the public, intuition, or flipping through a phone book—will help deter subversive action, especially from those who are not full-time dissidents.

Security has never been among the Internet's strong sides, and the proliferation of social media in the last decade has only made things worse. Even the most protected email service won't protect your password if there is a keylogger—software that can record and transmit your every keystroke—installed on your computer (or that slow and funky computer in a random Internet café that you once had to use). Nor does one need to break into your email to read some of it. Mounting and hiding a tiny, almost invisible digital camera behind your back is enough. Similarly, even secure, encrypted services like Skype will be of little consolation if a secret police operator occupies an apartment next to yours and sticks a parabolic microphone out the window. As long as most virtual activities are tied to physical infrastructure—keyboards, microphones, screens—no advances in encryption technology could eliminate all the risks and vulnerabilities.

But as security professionals attest, while it's possible to minimize the risks created by the infrastructure, it's much harder to discipline the users of a technology. Many sophisticated attacks originate by manipulating our trust networks, like sending us an email from a person we know or having us download files from trusted websites, as happened in the case of the Vietnamese activists. When we visit a website of an organization we trust, we do not expect to be hit with malware any more than we expect to be poisoned at a dinner party; we trust that the links we click on won't lead to sites that will turn our computers into mini-panopticons. Such trust has undoubtedly made the Internet an appealing place to do business or just waste so many hours of our lives. Few of us spend much time pondering the security settings on our favorite sites, especially if no sensitive data is divulged. But a low level of awareness is precisely what makes compromising the security of such sites so tempting, especially if these are niche sites catering to particular audiences. An attack can infect computers of all independent journalists, brave human rights defenders, or revisionist historians without triggering any suspicions from more computer-savvy user groups.

Poorly secured sites of specific communities thus enable the kind of attacks—many of which invariably result in more surveillance—that may not succeed were members of such communities targeted individually. This is what happened to Reporters Without Borders (RSF), a prominent international NGO defending freedom of expression, in July 2009, when someone inserted a malicious link into an email that RSF sent to its supporters. The link was placed next to the text of a 13,000-strong petition demanding the release of the documentary filmmaker Dhondup Wangchen from prison. Once clicked, it did lead to what looked like a genuine petition—so one would not suspect anything inappropriate—but the website also contained a security trap, infecting the computers of anyone who clicked on the malicious link. Alerted to the problem, RSF promptly removed the link, but it is difficult to estimate how many computers were compromised.

Even popular and much better-staffed organizations are not immune to embarrassing vulnerabilities that could cause damage to everyone in their social and professional circle. In early 2009 the website of the *New York Times,* which relies on banner ads provided by third parties, inadvertently served malware to some of its visitors. Such gaffes are poised to become even more widespread, as more and more websites incorporate a bevy of third-party services (e.g., Facebook's "like" button), surrendering full control over what kind of data flows through their site. When even the website of the *New York Times* feeds you viruses, there is little on the Internet you can safely surf on autopilot.

The Internet runs on trust, but its dependence on trust also opens up numerous vulnerabilities. Its effectiveness as a tool of carving out spaces of dissent and, in exceptional cases, even campaigning against authoritarian governments has to be judged on a much wider set of criteria than just the cost and ease of communications. It's quite obvious that in a world where there are no other uses for the Internet, email is a cheaper, more effective, and more secure alternative to the handwritten letter. But in a world like ours, where the Internet has many other functions, it would be a mistake to evaluate the practice of email in isolation from other online activities: browsing, chatting, typing, gaming, file sharing, and downloading and viewing porn. Each of these activities creates multiple vulnerabilities that alter the risk calculus.

It's important to avoid falling victim to Internet-centrism and focusing only on the intrinsic qualities of online tools at the expense of studying how those qualities are mitigated by the contexts in which the tools are used. Sending and receiving email on an Internet café's computer where the previous customer was downloading porn from illegal websites may not be a tremendous improvement over hand-delivering a typewritten letter. Yet this is the environment in which many activists in the developing world, short on money and equipment or simply hiding from the all-seeing eye of the secret police, are forced to work. Understanding the full gamut of risks and vulnerabilities that activists expose themselves to requires a bit more investigative work than simply comparing the terms of service that come with all newly created email accounts.

Why Databases Are Better Than Stasi Officers

Information may, indeed, be the oxygen of the modern age, as Ronald Reagan famously alleged, but it could be that peculiar type of oxygen that helps to keep dictators on life support. What reasonable dictator passes up an opportunity to learn more about his current or future enemies? Finding effective strategies to gather such information has always been a priority for authoritarian governments. Often such strategies were intrusive, such as placing bugs in dissidents' apartments and wiretapping their phone conversations, as happened in many countries of the Soviet bloc. But sometimes governments found more creative ways to do it, especially if they were simply trying to gauge public sentiment rather than peep inside the minds of particular dissidents.

The Greek military regime, for example, tried to keep track of everyone's reading habits by monitoring their choice of newspapers, thus quickly learning about their political leanings. The Greek generals would have loved the Internet. Today one could simply data-mine Amazon.com's wish lists—collections of books, films, and other items—that customers freely self-disclose. In 2006 the technology consultant Tom Owad conducted a quirky experiment: In less than a day he downloaded the wish lists of 260,000 Americans, used the publicly disclosed names and some limited contact information of Amazon's customers to find their full addresses, and then placed those with interesting book requests—like Orwell's *1984* or the Quran—on a map of the United States.

How do other old-school surveillance tactics score in the digital age? At first glance, it may seem they don't do so well. As a vast chunk of political communication has migrated online, there is little to be gained from bugging dissidents' apartments. Much of the digital information is swapped in silence, punctuated, perhaps, only by keystroke sounds; even the most advanced recording equipment cannot yet decipher those. Not surprisingly, analog bugs have long been replaced by their digital equivalent, making surveillance easier and less prone to error and misinterpretation; instead of recording the sounds of keyboard strokes, the secret police can now record the keyboard strokes themselves.

The Lives of Others, a 2006 Oscar-winning German drama, with its sharp portrayal of pervasive surveillance activities of the Stasi, GDR's secret police, helps to put things in perspective. Focusing on the meticulous work of a dedicated Stasi officer who has been assigned to snoop on the bugged apartment of a brave East German dissident, the film reveals just how costly surveillance used to be. Recording tape had to be bought, stored, and processed; bugs had to be installed one by one; Stasi officers had to spend days and nights on end glued to their headphones, waiting for their subjects to launch into an antigovernment tirade or inadvertently disclose other members of their network. And this line of work also took a heavy psychological toll on its practitioners: the Stasi anti-hero of the film, living alone and given to bouts of depression, patronizes prostitutes—apparently at the expense of his understanding employer.

As the Soviet Union began crumbling, a high-ranking KGB officer came forward with a detailed description of how much effort it took to bug an apartment:

> Three teams are usually required for that purpose: One team monitors the place where that citizen works; a second team monitors the place where the spouse works. Meanwhile, a third team enters the apartment and establishes observation posts one floor above and one floor below the apartment. About six people enter the apartment wearing soft shoes; they move aside a bookcase, for example, cut a square opening in the wallpaper, drill a hole in the wall, place the bug inside, and glue the wallpaper back. The artist on the team airbrushes the spot so carefully that one cannot notice any tampering. The furniture is replaced, the door is closed, and the wiretappers leave.

Given such elaborate preparations, the secret police had to discriminate and go only for well-known high-priority targets. The KGB may have been the most important institution of the Soviet regime, but its resources were still finite; they simply could not afford to bug everyone who looked suspicious. Despite such tremendous efforts, surveillance did not always work as planned. Even the toughest security officers—like the protagonist of the German film—had their soft spots and often developed feelings of empathy for those under surveillance, sometimes going so far as to tip them off about upcoming searches and arrests. The human factor could thus ruin months of diligent surveillance work.

The shift of communications into the digital realm solves many of the problems that plagued surveillance in the analog age. Digital surveillance is much cheaper: Storage space is infinite, equipment retails for next to nothing, and digital technology allows doing more with less. Moreover, there is no need to read every single word in an email to identify its most interesting parts; one can simply search for certain keywords—"democracy," "opposition," "human rights," or simply the names of the country's opposition leaders—and focus only on particular segments of the conversation. Digital bugs are also easier to conceal. While seasoned dissidents knew they constantly had to search their own apartments looking for the bug or, failing that, at least tighten their lips, knowing that the secret police was listening, this is rarely an option with digital surveillance. How do you know that someone else is reading your email?

To its credit, a few weeks after Google discovered that someone was trying to break into the email accounts of Chinese human rights dissidents, it began alerting users if someone else was also accessing their account from a different computer at that time. Few other email providers followed Google's lead—it would be seen as yet another unjustified expense—so this incident hardly put an end to the practice of secret police reading dissidents' email.

More important, the Internet has helped to tame the human factor, as partial exposure, based on snippets and keywords of highlighted text, makes it less likely for police officers to develop strong emotional bonding with their subjects. The larger-than-life personalities of fearless dissidents that melted the icy heart of the Stasi officer in *The Lives of Others* are barely visible to the Internet police, who see the subjects of surveillance reduced to one-dimensional, boring database entries. The old means of doing surveillance usually began with a target and only then searched for the crimes one could ascribe to it. Today, the situation is the reverse: Crimes—antigovernment slogans or suspicious connections to the West—are detected first, and their perpetrators are located later. It's hard to imagine Iranian Internet police developing sympathy for the people they investigate based on snippets of texts detected by the system, for they already know of their guilt and can always dig up more textual evidence if needed.

That technology helps to eliminate the indecision and frailty (and, more often than not, common sense and humanity) associated with human decision makers was not lost on the Nazis. Testifying at the

Nuremberg trials in 1946, Albert Speer, who served as Hitler's chief architect and later as the minister of armaments and war production, said that "earlier dictators during their work of leadership needed highly qualified assistants, even at the lowest level, men who could think and act independently. The totalitarian system in the period of modern technical development can dispense with them; the means of communication alone make it possible to mechanize the subordinate leadership." It's undoubtedly barbaric to be blaming Nazi atrocities on the evils of technology alone, but Speer had a point: The world is yet to meet a database that cried over its contents.

Tremendous cost savings introduced by digital surveillance technologies have also made it possible to shift surveillance personnel to more burning tasks. In a 2009 interview with *Financial Times,* a marketing manager for TRS Solutions, a Chinese data-mining firm that offers an Internet-monitoring service to the Chinese authorities, boasted that China's Internet police—thanks in part to the innovations developed by TRS Solutions—now only need one person where ten were required previously. But it's too early to celebrate; it's unlikely that the other nine were laid off. Most probably they were shifted to perform more analytical tasks, connecting the dots between hundreds of digital snippets gathered by automated computer systems. As the TRS manager pointed out, business is booming: "[The Chinese authorities have] many different demands—early warning, policy support, competitive spying between government departments. In the end, this will create a whole industry." Perhaps, this is not the kind of Internet-friendly industry celebrated by the proponents of wikinomics, who rarely acknowledge that, while the Internet has indeed helped to cut the unnecessary slack from many an institution, it has also inadvertently boosted the productivity of the secret police and their contractors in the private sector. A book on "wikiethics" is long overdue.

Say Hi. You're on Camera!

It's not just text that has become easier to search, organize, and act on; video footage is moving in that direction as well, thus paving the way for even more video surveillance. This explains why the Chinese government keeps installing video cameras in its most troubling cities. Not only do such cameras remind passersby about the panopticon they inhabit, they also supply the secret police with useful clues (in 2010 47,000 cameras were already scanning Urumqi, the capital of China's restive Xinjiang Province, and that number was projected to rise to 60,000 by the end of the year). Such revolution in video surveillance did not happen without some involvement from Western partners.

Researchers at the University of California at Los Angeles, funded in part by the Chinese government, have managed to build surveillance software that can automatically annotate and comment on what it sees, generating text files that can later be searched by humans, obviating the need to watch hours of video footage in search of one particular frame. (To make that possible, the researchers had to recruit twenty graduates of local art colleges in China to annotate and classify a library of more than two million images.) Such automation systems help surveillance to achieve the much needed scale, for as long as the content produced by surveillance cameras can be indexed and searched, one can continue installing new surveillance cameras.

But as the maddening pace of innovation in data analysis expands the range of what is possible, surveillance is poised to become more sophisticated as well, taking on many new features that only seemed like science fiction in the not-so-distant past. Digital surveillance is poised to get a significant boost as techniques of face-recognition improve and enter the consumer market. The face-recognition industry is so lucrative that even giants like Google can't resist getting into the game, feeling the growing pressure from smaller players like Face.com, a popular tool that allows users to find and automatically annotate unique faces that appear throughout their photo collections. In 2009 Face.com launched a Facebook application that first asks users to identify a Facebook friend of theirs in a photo and then proceeds to search the social networking site for other pictures in which that friend appears. By early 2010, the company boasted of scanning 9 billion pictures and identifying 52 million individuals. This is the kind of productivity that would make the KGB envious.

One obvious use of face-recognition technology would be to allow Iranian authorities to quickly learn the identity of the people photographed during street protests in Tehran. For why should the Iranian government embark on expensive investigations if they can get their computers to match the photos taken during the protests—many of them by the very activists appearing on them—with more casual photos uploaded on social networking profiles by the same activists? That said, governments and law-enforcement agencies had been using face-recognition technologies for a while before they became a commercially viable business. What is most likely to happen in the case of Iran is that widely accessible face-recognition technologies will empower various solo agents, socially and politically conservative cyber-vigilantes who do not work for the government but would like to help its cause. Just as hordes of loyal Thais surf the Web in search of websites criticizing the monarchy or hordes of pro-government Chinese are on the lookout for highly sensitive blog posts, hordes of hard-line Iranians will be checking photos from the antigovernment protests against those in massive commercial photo banks, populated by photos and names harvested from social networking sites, that are sure to pop up, not always legally, once face-recognition technology goes fully mainstream. The cyber-vigilantes may then continue stalking the dissidents, launch DDoS attacks against their blogs, or simply report them to authorities.

Search engines capable of finding photos that contain a given face anywhere on the Internet are not far on the horizon either. For example, SAPIR, an ambitious project funded by the European Union, seeks to create an audiovisual search engine that would first automatically analyze a photo, video, or sound recording; then extract certain features to identify it; and finally use these unique identifiers to search for similar content on the Web. An antigovernment chant recorded from the streets of Tehran may soon be broken down into individual voices, which in turn can then be compared to a universe of all possible voices that exist on amateur videos posted on YouTube.

Or consider Recognizr, the cutting-edge smartphone application developed by two Swedish software firms that allows anyone to point their mobile phone at a stranger and immediately query the Internet about what is known about this person (or, to be more exact, about this person's face). Its developers are the first to point to the tremendous privacy implications of their invention, promising that strict controls would eventually be built into the system. Nevertheless, it takes a leap of faith to believe that once the innovation genie is out of the bottle, no similar rogue applications would be available for purchase and download elsewhere.

How to Lose Face on Facebook

One gloomy day in 2009, the young Belarusian activist Pavel Lyashkovich learned the dangers of excessive social networking the hard way. A freshman at a public university in Minsk, he was unexpectedly called to the dean's office, where he was met by two suspicious-looking men who told him they worked for the KGB, one public organization that the Belarusian authorities decided not to rename even after the fall of communism (they're a brand-conscious bunch).

The KGB officers asked Pavel all sorts of detailed questions about his trips to Poland and Ukraine as well as his membership in various antigovernment movements.

Their extensive knowledge of the internal affairs of the Belarusian opposition—and particularly of Pavel's own involvement in them, something he didn't believe to be common knowledge—greatly surprised him. But then it all became clear, when the KGB duo loaded his page on vkontakte.ru, a popular Russian social networking site, pointing out that he was listed as a "friend" by a number of well-known oppositional activists. Shortly thereafter, the visitors offered Lyashkovich to sign an informal "cooperation agreement" with their organization. He declined—which may eventually cost him dearly, as many students sympathetic to the opposition and unwilling to cooperate with authorities have been expelled from universities in the past. We will never know how many other new suspects the KGB added to its list by browsing Lyashkovich's profile.

Belarus is not an isolated case, and other governments are quickly beginning to understand the immense intelligence value of information posted to social networking sites. Some even want to run their own sites, perhaps to save on surveillance costs. In May 2010, having banned Facebook and sensing the unmet and growing demands for social networking services among their population, Vietnam's Ministry of Information and Communications moved in to open their own social networking site, staffed with three hundred computer programmers, graphic designers, technicians, and editors. It is hard to say if it will become popular—with a name like GoOnline, it seems like a long shot—but from a government's perspective, it is even easier to spy on members of a social network once it knows all their passwords.

Democratic governments have also succumbed to such practices. The Indian police in the disputed territory of Kashmir, for example, are paying close attention to anything Kashmir-related that is posted on Facebook. On finding something suspicious, they call the users, ask about their activities, and order them to report to police stations. (This has prompted many activist users in Kashmir to start registering under false names, a practice that Facebook, keen not to dilute the quality of its superb user base with false entries, strongly discourages.)

Not all social networking is harmful, of course. Being part of a network carries many advantages. For example, it's much easier and cheaper to reach other members when such a need arises (e.g., before an upcoming protest). But membership in a network is something of a double-edged sword: Its usefulness can easily backfire if some segments get compromised and their relationships with other members become common knowledge. Before the advent of social media, it took a lot of effort for repressive governments to learn about the people dissidents are associated with. The secret police may have tracked one or two key contacts, but creating a comprehensive list—with names, photos, and contact information—was extremely expensive. In the past, the KGB resorted to torture to learn of connections between activists; today, they simply need to get on Facebook.

Unfortunately, there is still a widespread belief that authoritarian governments and their security services are too dumb and technophobic to go on social networking sites in search of such data. In his 2007 book *Children of Jihad* the U.S. State Department's Jared Cohen writes that "the Internet is a place where Iranian youth can operate freely, express themselves, and obtain information on their own terms. [They] can be anyone and say anything they want as they operate free from the grips of the police-state apparatus. . . . It is true that the government tries to monitor their online discussions and interactions, but this is a virtually impossible enterprise." This is simply factually wrong, as proven by the aftermath of the 2009 protests; for someone charged with developing effective Internet policy on Iran, Cohen is given to dangerously excessive cyber-utopianism. (One could only hope that it was not Cohen's Panglossian optimism that Condoleezza Rice, who hired him to work for the State Department's policy planning unit, was praising when she said that "Jared had insights into Iran that we [in the U.S. government] didn't have.") As it turns out, the Iranian authorities did spend a lot of time analyzing social networking sites in the aftermath of the elections and even used some of the information they gleaned to send warnings to Iranians in the diaspora. During the 2009 witch hunt trials in Iran, authorities used a dissident's membership in an academic mailing list run by Columbia University as proof that he was spying for Western powers.

Thus, even if an online social network is of minimum intelligence value, being friends with the wrong people provides evidence that can be used in court. Previously such information was hard to discover; often dissidents took extra efforts to conceal it. Belinda Cooper, an American activist who spent the late 1980s in GDR and was a member of several dissident environmental groups, writes that one of the rules practiced by the dissidents entering and leaving East Germany was to "never bring address books when going to the east (as border guards could and would photocopy them)." Today the situation has changed dramatically, as the lists of our friends on Facebook are available for anyone to see. Unfortunately, staying out of Facebook is not a reasonable option for most dissidents. They need to be present in these spaces to counter government propaganda, to raise awareness about their work in the West, to mobilize support for their causes among domestic audiences, and so forth. They may do so anonymously, of course, but anonymity also makes their involvement far less effective. Sakharov's advocacy would have been far less successful if he hadn't practiced it openly.

Numerous academic studies confirm that every time we share personal data on a social networking site, we make it more likely that someone might use it to predict what we are like, and knowing what we are like is a good first step toward controlling our behavior. A 2009 study by researchers at MIT has shown that it is possible to predict—with a striking degree of accuracy—the sexual orientation of Facebook users by analyzing their online friends. This is hardly good news for those in regions like the Middle East, where homosexuality still carries a heavy social stigma.

Another 2009 study conducted by researchers at the University of Cambridge, whose report is titled "Eight Friends Are Enough," found that based on the limited information that Facebook discloses to search engines like Google, it is possible to make accurate inferences about information that is not being disclosed.

Many of the functions that make social networking sites so easy to use—for example, to find one's friends who are already members of the site—also make it easy to trace identities behind emails or even trace users' activities across various other sites. Most of us know how easy it is to check whether our friends have already signed up for particular social networking sites simply by granting Facebook, Twitter, or LinkedIn temporary access to our email address book, so that those sites can automatically check the email addresses of our contacts against their lists of existing users. If five of our email buddies are already Twitter users, Twitter can let us know. So far, so good. The problem is that one can do the same operation with one's enemies as well. Email addresses can be added to address books manually, without ever having to email that person. Thus, just by knowing a person's email address, it might be possible to find her accounts on all social networking sites, even if she doesn't use those sites under her real name.

A 2010 study by Eurecom, a French research institute, sought to investigate the security vulnerabilities that such ease of use creates for the user. First, the researchers found 10.4 million email addresses on the Web; then they imported them into their address books; and, finally, they developed a simple script to automatically check with each of the popular social networking sites whether it had any users corresponding to those emails. As a result, they identified more than 876,941 emails linked to 1,228,644 profiles, with 199,161 emails having accounts on at least two sites, 55,660 on three, and so forth (11 people had their email accounts linked to seven social networking sites at once).

As was to be expected, some users who had accounts on multiple social networking sites provided different details to each (for example about their sexual orientation, location, or age). It's highly probable that quite a few of the people under investigation didn't want anyone to link the kind of frivolities they post to Twitter with their line of work, and yet researchers found at least 8,802 users who had accounts on both LinkedIn, a social network for professionals, and Twitter. If someone in that pool listed, say, "U.S. Department of Defense" as their employer on their LinkedIn profile, one could check what that person was tweeting about, even if the tweeting was done under a nickname.

Therefore, as long as social networking accounts are tied to one email address, it's also remarkably easy to tie them to a particular person, learn that person's name, and see what kind of hidden indiscretions that person may be engaging in, offline or online. The researchers, for example, found the profile of a married professor in his fifties who was also remarkably active on various dating sites. Similarly, activists who upload sensitive videos to YouTube thinking that no one could guess their real names from their usernames may be under much greater risks if they use the same email address to access Facebook and the secret police learns what that email address is.

Once alerted to such vulnerabilities, many social networking sites slightly tweaked their operating procedures, making it hard to do such checks in bulk. Nevertheless, it's still possible to find multiple online identities for individual emails through manual checking. This is not the kind of feature that is going to disappear soon, if only because it allows social networking sites to expand their user base.

Corporations are already taking advantage of the increasingly social nature of the Web. Hotels now use locations, dates, and usernames that appear on sites like TripAdvisor or Yelp to triangulate a guest's identity. If they find a likely match and the review happens to be positive, the review is added to a hotel's guest preference records. If it's negative, the travelers might be given a voucher to compensate for the inconvenience or, in the worst scenario, to be marked as "problem guests." Barry Hurd, the CEO of Seattle-based 123 Social Media, a reputation management company that works with more than five hundred hotels, believes that "technology is evolving so fast that in the future, every hotel representative could have a toolbar on his or her computer that reveals everything about a guest at the click of a mouse—every review, guest preference and even the likelihood that you'll be positively or negatively inclined toward your stay."

Of course, hotels are not authoritarian governments—they won't imprison guests in their rooms for expressing dissenting views—but if they can learn the real identities behind imaginary online nicknames, so can the secret police. Moreover, the corporate quest for de-anonymizing user identities can soon fuel a market in tools that can automate the process, and those tools can then be easily used in more ominous contexts. Intelligence agencies in the United States have already profited from data-gathering technology created on Wall Street. TextMiner, one such platform developed by Exegy, a firm that works with both intelligence agencies and Wall Street banks, can search through flight manifests, shipping schedules, and phone records as well as patterns that might form Social Security numbers or email accounts. "What was taking this one particular agency one hour to do, they can now do in one second," says Ron Indeck, Exegy's chief technology officer, in a phrase that sounds remarkably similar to the glee of the Chinese contractors at TRS Solutions. Thus, an entire year's worth of news articles from one organization can be searched and organized in "a couple of seconds." The private sector will surely continue churning out innovations that can benefit secret police everywhere. Without finding ways to block the transfer of such technologies to authoritarian states or, even more important, the kind of limits that should be imposed on such technologies everywhere, the West is indirectly abetting the work of the secret police in China and Iran.

But even in the absence of such tools, creative hacks will do the job just fine. A 2010 collaborative project between researchers at the Vienna University of Technology, the University of California at Santa Barbara, and Eurecom found an interesting way of de-anonymizing users of Xing, a popular German social networking site akin to Facebook and LinkedIn. Since most of us belong to a number of different social networking groups that vary according to our passions, life history, and lifestyle—for example, Save the Earth, Feed the Children of Africa, Alumni of the Best University in the World, Vegetarians of the World Unite—the probability that you and your friends belong to exactly the same groups is small (having attended the same liberal arts college in New England, your best friend may also want to save the earth and feed the children of Africa but also love Texas barbecue ribs).

Social networking sites do not usually hide lists of group members from nonmembers, so as not to erect too many communication barriers. It is thus possible to produce a nearly unique identifier, a "group fingerprint"—think of this as a list of all Facebook groups that a given user belongs to—for each of us. And the most obvious place to look for a matching fingerprint would be in our web browsers' history, for this is where a record of all the groups—and, of course, of all other websites we visit—is kept. All it takes to steal our browser history is to have us click on a malicious link, like the one mysteriously added to RSF's email petition, and everything we have been browsing in the last few days will no longer be private knowledge.

According to the 2010 report, producing a matching "group fingerprint" required the checking of 92,000 URLs, which took less than a minute. The researchers managed to correctly guess the identity of their target 42 percent of the time. In other words, if someone knows your Web history and you happen to be an avid user of social networking sites, she has a good chance of deducing your name. Soon, the secret police will just be able to look at the log from your favorite Internet café and learn who you are, even without asking for a copy of your passport (although that latter option is also increasingly common in authoritarian governments).

It's hardly surprising that the secret police in authoritarian regimes are excited about exploiting such vulnerabilities to fill in gaps in their databases. They may, for example, know email addresses of government opponents but not their identities. To learn their names, they could send the opponents fake emails containing malicious links that aim to steal their browsers' histories. In just a few minutes, they'll be able to attach names (as well as photos, contact details, and information about related connections) to their rather sparse database entries. Another problem is that social networking sites like Facebook don't thoroughly screen external developers—those who work on all those online games, quizzes, and applications—for trustworthiness. (Until very recently, they also did not impose clear limits on how much user data such applications could have access to, regardless of their actual needs.) This means, in essence, that a smart authoritarian regime can just put together a funny quiz about Hollywood movies and use it to gather sensitive information about its opponents. This is a nightmarish scenario for activists who struggle to keep their connections hidden from authorities; obviously, if the government knows all the Facebook friends of its fiercest political opponents, it would be silly not to pay close attention to their online activities, too, as there is always a good chance they also pose a threat.

Nor does it help that in their ill-conceived quest for innovation, technology companies utterly disregard the contexts in which many of their users operate, while significantly underestimating the consequences of getting things wrong. In early 2010, when Google launched Google Buzz, its Twitter-like service, they did not take appropriate care in protecting the identities of many of their users, disclosing their contact lists in the erroneous belief that no one would mind such intrusions into their privacy (even Andrew McLaughlin, Google's former senior executive and the deputy chief technology officer in the Obama administration, was trapped in the Buzz trap, as many of his former Google colleagues appeared in his contact list). Though Google executives downplayed the significance of the accident by claiming that no one got seriously hurt in the debacle, in truth we don't know how many new names and connections were added to the KGB's databases as a result. The real costs of Google's misjudgment cannot be immediately calculated.

Think, Search, Cough

Every time we post a greeting to our friend's Facebook wall, Google the name of our favorite celebrity, or leave a disapproving comment on the website of our favorite newspaper, we leave a public trail somewhere on the Internet. Many of these trails, like the comment on the newspaper's site, are visible to everyone. Some, like our Google searches, are only visible to us (and, of course, Google). Most, like that odd comment on the Facebook wall, fall somewhere in between.

Fortunately, we are not alone on the Internet—at least one billion other users are also blogging, Googling, Facebooking, and tweeting—and most of our information is simply lost in the endless ocean of digital ephemera produced by others. This is what privacy scholars call "security by obscurity." In most cases, obscurity still works, even though there are more and more exceptions to this rule. Ask anyone who has difficulty finding a job or renting an apartment because something embarrassing about him or her appears in Google searches or on Facebook. Nevertheless, aggregating these tiny digital trails into one big data set—sometimes across entire populations—could produce illuminating insights into human behavior, point to new trends, and help predict public reaction to particular political or social developments. Marketing and advertising companies understood the power of information a long time ago. The more they know about demographics, consumer habits, and preferences of particular customer types, the more they can tailor their product offerings, and the more they can make in sales as a result.

The digital world is no different. The history of our Internet search says more about our information habits than our patron files in the local library. The ability to identify and glean "intent" from a mere Internet search, matching advertisers with customers looking for their offerings, has allowed Google to turn the advertising business on its head. Thus, in addition to running the world's most successful advertising agency, Google also runs the most powerful marketing intelligence firm. This is because Google knows how to relate Internet searches to demographics and other searching and purchasing decisions of its customers (e.g., what percent of New Yorkers who searched for "digital camera" in the past twelve months ended up searching for "deals on iPhones").

But we're not just looking for better iPods and new deals on plasma TVs. We are also seeking information about people and places in the news ("has Michael Jackson died?"), about broader cultural trends ("what are the best novels of the decade?"), and, of course, about solving problems—mostly trivial but some important—that constantly pop up in our lives ("how to repair a broken washing machine").

There are many seasonal variations to how often we search for particular items (searches for "stuffed turkey" predictably increase before Thanksgiving), but the frequency of queries for most items is usually fairly consistent. Thus, whenever there is a sudden spike in the number of Google queries for a given term, it probably indicates that something extraordinary has just happened; the likelihood is even higher if the search spike is limited to a particular geographic area only.

For example, when an unusually high number of Internet users in Mexico began Googling terms like "flu" and "cold" in mid-April 2009, it signaled the outbreak of swine flu. In fact, Google Flu Trends, a dedicated Google service built especially for the purpose of tracking how often people search for flu-related items, identified the spike on April 20, before the swine flu became a cause célèbre with many in the media. And even though several scientific studies by health researchers found that Google's data is not always as accurate as other ways of tracking the spread of influenza, even they acknowledged how cheap and quick Google's system is. Besides, in fields that are not as data-intensive as disease control, Google does a much better job than the alternatives—if those exist at all.

Search engines have inadvertently become extremely powerful players in the business of gathering intelligence and predicting the future. The temptation—which Google executives, to their credit, have resisted so far—is to monetize the vast quantity of this trends-related information beyond just ad sales.

Technically, Google does know how often Russian Internet users search for the words "bribes," "opposition," and "corruption"; it even knows how such queries are distributed geographically and what else such potential troublemakers are searching for. It does not take a Nostradamus to interpret a sudden spike of Internet searches for words like "cars," "import," "protests," and "Vladivostok" as a sign of growing social tensions over increases in car tariffs brewing in Vladivostok, Russia's major outpost in the Far East.

This is the kind of data that Russian secret services would literally kill for. Such knowledge may, of course, make authoritarianism more responsive and inject at least a modicum of democracy into the process. But it's also possible that governments would use this knowledge to crack down on dissenters in a more effective and timely manner.

Internet search engines offer an excellent way to harness the curiosity of the crowds to inform the authorities of impending threats. Monitoring an Internet search could produce even more valuable intelligence than monitoring Internet speech, because speech is usually directed at somebody and is full of innuendo, while an Internet search is a simple and neutral conversation between the user and the search engine.

The intelligence value of search engines is not lost on the Internet gurus consulting authoritarian governments. In March 2010, speaking about the Kremlin's ambitions to establish its own search engine, Igor Ashmanov, one of the pioneers of the Russian Internet and someone who had consulted for the Kremlin about their national search plan in the past, was direct: "Whoever dominates the search market in the country knows what people are searching for; they know the stream of search queries. This is completely unique information, which one can't get anywhere else." If one assumes that authoritarian governments usually fall by surprise—if they are not surprised, they are probably committing suicide (e.g., the case of the Soviet Union)—then we also have to assume that, given how much data on the Internet can be harvested, analyzed, and investigated, surprises may become rarer.

But even if the governments' attempts to control—directly or indirectly—the world of Internet search would not bring immediate results, the Internet could boost their intelligence-gathering apparatus in other ways. The advent of social media has made most Internet users increasingly comfortable with the idea of sharing their thoughts and deeds with the world at large. It may not seem obvious, but trolling through all those blog posts, Twitter updates, photos, and videos posted to Facebook and YouTube could yield quite a lot of useful information for intelligence services—and not just about individual habits, as in the Belarusian KGB case, but also about broad social trends and the public mood as a whole. Analyzing social networks could offer even better insights than monitoring online searches, as one could correlate information coming from particular individuals (whether it's opinions or facts) in the light of what else could be known about these individuals from their social networking profile (how often they travel, what kind of online groups or causes they embrace, what movies they like, who else is in their network, etc.).

An authoritarian government, for example, may pay special attention to the opinions of those who are between twenty and thirty-five years old, frequently travel abroad, and have advanced degrees. One simply needs to spend some time browsing relevant Facebook groups (e.g., "Harvard class of 1998" or "I love traveling in the Middle East") to zero in on the right characters. In a sense, the world of social networking obviates the need for focus groups; finding smart ways to cluster existing online groups and opinions could be more effective. And they don't have to collect this data on their own. Plenty of private companies are already collecting data—mostly for marketing purposes—that governments, both authoritarian and democratic ones, would find extremely useful. Thus, while the KGB may no longer exist in 2020, its functions may still be performed by a smattering of private companies specializing in one particular aspect of information work.

Today governments can learn quite a lot about the prospects of political unrest in a particular country simply by paying particular attention to the most popular adjectives used by the digerati. Are they "happy" or "concerned"? Do they feel "threatened" or "empowered"?

What if one controls for religion? Do self-professed secular bloggers feel more satisfied than the religious ones?

Just imagine how useful it might be for the Iranian government to track how often Iranians use the word "democracy" in their public online conversations and how such mentions are spread across the country. (For example, are there any regions of Iran that are more democratically inclined and unhappy with the current regime than others?)

If proper controls for statistical bias are in place, such technology is often superior to opinion polls, which take time to develop and, when done in authoritarian countries, always carry the risk of people misrepresenting their views to avoid punishment. Such aggregated information may not be fully representative of the entire population, but it helps to keep the tab on the most troublesome groups. Thus, the fact that authoritarian governments can now learn more about the public mood in real time may only add to their longevity. They are less likely to misjudge the public reaction.

What's worse is that social media activity is not always a bad proxy for judging the relative importance of antigovernment activists. If tweets of a particular user are retweeted more often than average, it's a good idea for the government to start watching that individual closely and learn more about his or her social network. The viral culture of social media may at least indirectly help solve the problem of information overload that has affected censorship as well. It's the "online marketplace of ideas" that tells secret police whom to watch. From the perspective of the secret police, people who are unpopular probably don't even deserve to be censored; left to their own devices and nearly zero readers, they will run out of blogging energy in a month or so.

The Myth of an Overprotected Activist

Despite the terrifying efficiencies in the practice of surveillance that were introduced by digital technology, not all is lost. It would be disingenuous to suggest that the digital realm has nothing in store for dissidents; it has greatly enhanced many of their activities as well. One great intellectual challenge facing any scholar of today's Internet is being able to see the risks inherent in new technologies while not discarding the numerous security-enhancing opportunities that they offer. The only way to come up with a satisfying answer to the question of whether the Internet has eroded or strengthened the surveillance and control apparatus of authoritarian governments is to examine all major technologies one by one, in their specific contexts.

But first it may help to examine the ways in which the Internet has helped dissidents to conceal antigovernment activities. First, sensitive data can now be encrypted on the cheap, adding an extra level of protection to conversations between dissidents. Even though decryption is possible, it can eat a lot of government resources. This is particularly true when it comes to voice communications. While it was relatively easy to bug a phone line, this is not such an easy option with voice-over-the-Internet technology like Skype. (The inability to eavesdrop on Skype conversations bothers Western governments, too: In early 2009 the U.S. National Security Agency was reported to have offered a sizeable cash bounty to anyone who could help them break Skype's encrypted communications; to date no winners have been announced.)

Second, there is so much data being produced online that authorities cannot possibly process and analyze all of it. Comparable estimates for the developing world are lacking, but according to a 2009 study by researchers at the University of California at San Diego, by 2008 the information consumption of an average American reached thirty-four gigabytes of data per day, an increase of 350 percent compared to 1980. The secret police have no choice but to discriminate; otherwise, they may develop a severe case of attention deficit disorder, getting bogged down in reading millions of blogs and Twitter updates and failing to see the big picture. Thanks to this data deluge, it may take a few months before authorities discover the new hideout of activists, who thus gain a few months of unsupervised online collaboration. The authorities are much better informed about the parameters of the haystack, but the needle is still quite hard to find.

Third, technologies like Tor now make it possible to better protect one's privacy while surfing the Internet. A popular tool that was initially funded by the U.S. Navy but eventually became a successful independent project, Tor allows users to hide what it is they are browsing by first connecting to a random "proxy" node on the volunteer Tor network and then using that node's Internet connection to connect to the desired website. Interestingly, as users of the Saudi site Tomaar found out, tools like Tor also help to circumvent government filtering of the Internet, for, from the government's perspective, the user is not browsing banned websites but is simply connecting to some unknown computer. This is why once the Iranian government found out the proxies used by its opponents during the 2009 protests, many of them publicized by unsuspecting Westerners on Twitter, it immediately began blocking access to them.

But Tor's primary function remains guaranteeing its users' anonymity. Think of this as surfing the Internet using an anonymous network of helpers who fetch all the websites you need and thus ensure that you yourself are not directly exposed. As long as the government doesn't know these helpers by name, the helpers don't know each other, and you frequent enough other networks not to attract attention to the helpers, you can get away with browsing whatever you want.

But how many activists actually bother to read the fine print that is invariably attached to all modern technologies? Most probably ignore it. If the Soviet dissidents had to memorize the manuals to their smuggled photocopiers before distributing any samizdat, their output might have been considerably less impressive. And a lot of the tools are easy to misunderstand. Many users, including those in the most secretive government outfits, mistakenly believe that Tor, for example, is more secure than it actually is. Swedish researcher Dan Egerstad set up five Tor nodes of his own—that is, he became one of the final stage helpers—to learn more about data that passed through them. (The "helper" who finds herself as the final node on the network—that is, it helps to gain access to the desired target site rather than simply redirect the request to another "helper"—can see what websites it is actually "helping" to access, even though it won't know who is trying to access them.) Egerstad, who was arrested as a result of his little scholarly experiment, found that 95 percent of the traffic that passed through his experimental Tor connections—including government documents, diplomatic memos, and intelligence estimates—was not encrypted. Think of intercepting an envelope that doesn't have a return address. Would you be able to guess who wrote it? Sure, if you look inside: The letterhead may tell you everything you need to know. TOR is excellent at removing the sender's address from the envelope, but it doesn't destroy the letterhead, let alone the rest of the letter. There are, of course, plenty of other encryption technologies that can do this, but Tor is simply not one of them. That so many users exchanging sensitive information online—including activists and dissidents—do not have a firm understanding of the technologies they use is cause for serious concern. Eventually it puts them at completely unnecessary and easily avoidable risk.

Besides, even complete mastery of technology is often not enough. Your security is only as good as that of the computer you are working on; the more people have access to it, the more likely it is that someone could turn your computer into a spying machine. Given that a lot of Internet activism takes place on public computers, security compromises abound. For many antigovernment activists, cybercafés have become the new (and often the only) offices, as authorities keep a close eye on their home and office Internet connections. However, few Internet cafés allow their patrons to install new software or even use browsers other than Internet Explorer, which puts most innovative tools for secure communication out of easy reach.

Rainy Days of Cloud Computing

Some observers see many security-enhancing benefits to the Internet. For example, dissidents and NGOs can now use multifunctional online working environments to execute all their work remotely—"in the cloud"—without having to install any software or even store any data on poorly protected computers. All one needs is a secure browser and an Internet connection; there's no need to download any files or carry a portable copy of your favorite word processor on a USB thumb drive.

"Cloud activism" may, indeed, seem like something of a godsend, an ideal solution to data security concerns faced by many NGOs and activists. Take the case of Memorial, a brave Russian NGO that has gained worldwide recognition for its unyielding commitment to the documenting of human rights abuses and crimes committed in the country, from Stalin's rule to the more recent wars in Chechnya.

On November 4, 2008, only a day before an edgy conference on Stalin's role in modern Russia co-organized by Memorial, the Russian police raided its offices in Saint Petersburg and confiscated twelve hard drives containing the entire digital archives of atrocities under Stalin, including hours of audio histories and video evidence of mass graves. It was an institutional disaster. Not only did Memorial lose possession of (even if temporarily) twenty years of important work, but Russian authorities were supplied with potentially damning evidence against the organization. Given that historical memory—especially of the Stalin period—is a sensitive issue in Russia, finding fault with Memorial, which happens to be a staunch critic of the Kremlin, wouldn't be so hard. Russian police are notorious for finding fault with the most innocuous of documents or, worse, software and operating systems. (Quite a few Russian NGOs use illegal software in their offices, often without even realizing it until it is too late; on more than one occasion, the war on pirated software, which the West expects Moscow to fight with all its vigor, has been a good excuse to exert more pressure on dissenting NGOs.)

Fortunately, the courts concluded that the search had been conducted in violation of legal due process, and Memorial's hard drives were returned in May 2009. Nevertheless, the fact that authorities had simply walked in and confiscated twenty years of work posed a lot of questions about how activists might make digital data more secure.

Fans of "cloud activism" would point out that one way to avoid disasters like Memorial's is to shift all data into the cloud, away from local hard drives and onto the Internet, thus making it impossible for the authorities to confiscate anything. To get access to such documents, authorities would need a password, which, in most countries, they would not be able to obtain without a court order. (Of course, this would not work in countries that have absolutely no respect for the rule of law; one can learn the password by torturing the system administrator without having to go through the courts.)

The possibility of using online word-processing services like Google Docs and dumping all important data on the Internet may, indeed, seem like an improvement over storing data on easily damaged, insecure hard drives lying around NGOs' dusty offices. After all, the data could be stored on a remote server somewhere in California or Iowa, completely out of immediate reach of authoritarian governments, if only because it ensures that the latter cannot legally and physically get to the services storing it (or not immediately, at any rate).

While there is much to admire about this new cloud-based model, it also comes with tremendous costs, which could sometimes outweigh the benefits. One major shortcoming of producing and accessing documents in the cloud is that it requires a constant transmission of data between a computer and a server where the information is stored. This transmission is often done "in the open" (without proper encryption), which creates numerous security compromises.

Until very recently, many of Google's online offerings—including such popular services as Google Docs and Google Calendar—did not offer encryption as the default option. This meant that users connecting to Google Docs through, say, insecure Wi-Fi networks were playing with fire: virtually anyone could see what they were sending to Google's servers. Fortunately, the company altered its encryption policy after several high-profile security experts wrote a letter to its CEO, where they highlighted the unnecessary and easily avoidable risks to Google users. But Google is not the only player in this space—and where Google has the resources to spend on extra encryption, others may not. Making encryption the default setting may slow down the service for other users and impose new costs on the company's operations. Such improvements are not completely out of the question, however. A strong argument can be made—hopefully, by lawmakers on both sides of the Atlantic—that forcing Internet companies to enhance the security of their services makes a lot of sense from the perspective of consumer protection regulation. Instead of giving such companies a free pass because they are now the key players in the fight for Internet freedom, Western governments should continue looking for ways in which their services could be made extremely secure, for anything less than that would, in the long run, endanger too many people.

But other insecurities abound, too. The fact that many activists and NGOs now conduct all their business activities out of a single online system, most commonly Google—with calendar, email, documents, and budgets all easily available from just one account—means that should their password be compromised, they would lose control over all of their online activities. Running all those operations on their own laptops was not much safer, but at least a laptop could be locked in a safe. The centralization of information under one roof—as often happens in the case of Google—can do wonders from the perspective of productivity, but from the perspective of security it often only increases the risks.

On Mobile Phones That Limit Your Mobility

Much like cloud computing, the mobile phone is another activist tool that has not been subjected to thorough security analysis. While it has been rightly heralded as the key tool for organizing, especially in countries where access to the Internet and computers is prohibitively expensive, little has been said about the risks inherent to most "mobile activism."

The advantages of such activism are undeniable. Unlike blogging and tweeting, which require an Internet connection, text messaging is cheap and ubiquitous, and it doesn't require much training. Protesters using mobile phones to organize public rallies have become the true darlings of the international media. Protesters in the Philippines, Indonesia, and Ukraine have all taken advantage of mobile technology to organize and challenge their governments. This technology is not without its shortcomings and vulnerabilities, however.

First and foremost, authorities can shut down mobile networks whenever they find it politically expedient. And they do not have to cut off the entire country; it's possible to disconnect particular geographic regions or even parts of the city. For example, during the unsuccessful color revolution in Belarus in 2006, the authorities turned off mobile coverage in the public square where protesters were gathering, curbing their ability to communicate with each other and the outside world (the authorities claimed that there were simply too many people using mobiles on the square and the mobile networks couldn't cope with the overload). The Moldovan authorities made a similar move in spring 2009, when they turned off mobile networks in the central square of Chisinau, Moldova's capital, thus greatly hampering the communication capacity of those leading the local edition of the Twitter revolution. Such shutdowns can also be on a larger, national scale and last longer. In 2007 the government of Cambodia declared a "tranquility period," during which all three mobile operators agreed to turn off text messaging for two days (one of the official explanations was that it would help keep voters from being flooded with campaign messages).

Many authorities have mastered the art of keyword filtering, whereby text messages containing certain words are never delivered to their intended recipients. Or they may be delivered, but the authorities will take every step to monitor or punish their authors. In 2009 police in Azerbaijan reprimanded forty-three people who voted for an Armenian performer (Armenia and Azerbaijan are at war over the disputed Nagorno-Karabach territory) in the popular Eurovision contest, summoning some of them to police headquarters, where they were accused of undermining national security, and forced to write official explanations. The votes were cast by SMS. In January 2010, *China Daily,* China's official English-language newspaper, reported that mobile phone companies in Beijing and Shanghai began suspending services to cellphone users who were found to have sent messages with "illegal or unhealthy" content, which is the Chinese government's favorite euphemism for "smut."

This means that China's mobile operators would now be comparing all text messages sent by their users to a list of banned words and blocking users who send messages containing banned words. That's a lot of messages to go through: China Mobile, one of China's biggest mobile operators, processes 1.6 billion text messages per day. Even though the campaign officially claims to be fighting pornography, similar technology can be easily used to prevent the distribution of text messages on any topic; it all depends on the list of banned words. Not surprisingly, this list of "unhealthy words" comes from China's police. But there is also plenty of traffic in the other direction—that is, from companies to the state. Wang Jianzhou, China Mobile's CEO, stunned the attendees of the World Economic Forum in Davos in 2008 by claiming that his company provides data on its users to the government whenever the government demands it.

What's worse, Western companies are always happy to provide authoritarian governments with technology that can make filtering of text messages easier. In early 2010, as American senators were busy praising Google for withdrawing from China, another American technology giant, IBM, struck a deal with China Mobile to provide it with technology for tracking social networks (of the human, not virtual variety) and individuals' messaging habits: who sends what messages to whom and to how many people. (IBM, of course, was quick to point out that such technology is meant for helping Chinese mobile operators cut down on spam, but none can vouch that the same operators won't use it to curb political speech.)

Any technologies based on keyword filtering can, of course, be easily tricked. One can deliberately misspell or even substitute most sensitive words in a text message to fool the censors. But even if activists resort to misspelling certain words or using metaphors, governments could still make the most popular of such messages disappear. In fact, it's not the actual content of the messages that worries the government—no one has yet expressed cogent government criticism in a hundred forty characters or less—but the fact that such messages could go viral and be seen by millions of people. Regardless of the content being shared, such viral dissemination of information makes authoritarian governments feel extremely uneasy, as it testifies to how much their grasp on information has been eroded. In the most extreme cases, they won't hesitate to use the nuclear option and block most popular messages, without paying much attention to their content.

What is even more dangerous about using mobile phones for activism is that they allow others to identify the exact location of their owners. Mobile phones have to connect to local base stations; once a user has connected to three bases, it is possible to triangulate the person's position. In an online demonstration to its current and potential customers, ThorpeGlen, a U.K.-based firm, boasts that it can track "a specific target through ALL his electronic communications. . . . We can detect change of SIM and change of handset after identifying one suspect. . . . We can even detect that profile again even if the phone AND SIM are changed." This means that once you've used a cellphone, you are trapped. To clinch their marketing pitch, ThorpeGlen attached an online map of Indonesia that depicted the movements of numerous dots—millions of Indonesians with their cellphones; it allowed a viewer to zoom in on any particular sector. But it is hardly the only company offering such services; more and more start-ups cater to the vibrant consumer market in cellphone surveillance. For just $99.97 a year, Americans can load a little program called MobileSpy onto someone's cellphone and track that phone's location whenever they want.

Monitoring the geographic location of phone owners may enable the government to guess where big public actions might be happening next. For example, if the owners of the hundred most dangerous cellphone numbers are all seen heading to a particular public square, there is a good chance that an antigovernment demonstration will soon ensue. Furthermore, mobile companies have strong economic incentives to improve their location-identification technology, as it would allow them to sell geographically targeted advertising, such as prompts to check out the café next door. If anything, determining a person's location by tracing his or her mobile phone is poised to get easier in the future. While ThorpeGlen markets its services to law enforcement and intelligence firms in the West, it's not clear if any restrictions would prohibit the export of such technology elsewhere.

Many activists are, of course, aware of such vulnerabilities and are doing their best to avoid easy detection; however, their most favorite loopholes may soon be closed. One way to stay off the grid has been to buy special, unbranded models of mobile phones that do not carry unique identifiers present in most phones, which could make such devices virtually untraceable. Such models, however, also appeal to terrorists, so it's hardly surprising that governments have started outlawing them (for example, in the wake of 2008 attacks in Mumbai, India banned the export of such phones from China). The frequent use of new technologies by terrorists, criminals, and other extreme elements presents a constant challenge to Western governments who would like to both empower democratic activists and disempower many of the sinister nonstate groups that are undermining the process of democratization.

Another favorite low-tech solution, disposable prepaid SIM cards, which allow activists to change their phone numbers on a daily basis, may not stay around for much longer either, as buying them is becoming more difficult in many parts of the developing world. Russia and Belarus, for example, require retailers to obtain a copy of the customer's passport when someone buys a prepaid card, which essentially eliminates the desired anonymity. In early 2010 Nigeria passed a similar law, and other African states are expected to follow. Since American policymakers fret about Al-Qaeda jihadists using prepaid SIM cards to coordinate terrorist acts, it's quite likely that similar measures will soon pass in the United States as well. In 2010, with the entire country abuzz with the Times Square terror threat, FBI Director Robert Mueller endorsed anti-terrorism legislation that would require prepaid cellphone sellers to keep records of buyers' identities.

As useful as mobile technology could be for countering the power of authoritarian states, it comes with numerous limitations. This is not to say that activists should not be harnessing its communications power. They should, but only after fully familiarizing themselves with all the risks involved in the process.

As the Web becomes more social, we are poised to share more data about ourselves, often forgetting about the risks involved. Most disturbingly, we do so voluntarily, not least because we often find such sharing beneficial. Thus, sharing our geographical location may alert our friends to our whereabouts and facilitate a meeting that may not have happened otherwise. What we often overlook is that by saying where we *are*, we are also saying where we are *not*. Obviously, this is a boon for burglars; privacy activists even set up a dedicated site provocatively called "Please Rob Me" to raise public awareness about such risks. Such a wealth of data is also of great value to authoritarian states. Today's digitized, nimble, and highly social surveillance has little in common with the methods practiced by Stasi and KGB in 1989. The fact that there are more ways to produce and disseminate data has not overloaded the censorship apparatus, which has simply adapted to this new age by profiting from the same techniques—customization, decentralization, and smart aggregation—that have propelled the growth of the Internet. The ability to speak and make connections comes with costs, and those costs may not always be worth the benefits.

Denying that greater information flows, combined with advanced technologies like face or voice recognition, can result in the overall strengthening of authoritarian regimes is a dangerous path to take, if only because it numbs us to potential regulatory interventions and the need to rein in our own Western corporate excesses. It's not a given that IBM should be selling SMS-filtering technology to authoritarian states; that services like Google Buzz should be launched with minimum respect for the privacy of its users; that researchers at public universities like the University of California should be accepting funding from the Chinese government to work on better video surveillance technology; or that Facebook should be abdicating their responsibility to thoroughly screen developers of its third-party applications. All of these developments are the result of either excessive utopianism, unwillingness to investigate how technology is being used in non-Western contexts, or unquenchable thirst for innovation with complete disregard for its political consequences. While the Internet by itself may not be liberating those living in authoritarian states, Western governments should not be making it easier to use in suppressing dissent.