



INSIDE THE BATTLE FOR CYBERSPACE

.....

RONALD J. DEIBERT

2.

Filters and Chokepoints

“I have no idea what the Internet is!”

— Hayastan Shakarian, aged seventy-five

On March 28, 2011, the Internet went down in Georgia. For nearly twelve hours citizens had no access to Twitter, Facebook, their favourite YouTube videos, or their primary sources of news and online information. They could not access their online bank accounts or send emails. An information darkness had descended on the Eurasian country. The culprit? A nasty computer virus? Another Russian invasion? The latter would not be out of the question. Three years earlier, Georgia's Internet was brought to a halt as Russian ground troops invaded the territorial enclave of South Ossetia, the country's most contested region. Acting in support of the Motherland, scores of patriotic Russian hackers bombarded the Georgian Internet with a massive DDOS attack. It overwhelmed Georgian computers, including the government's websites and the country's banking and 911 systems.

As it turned out, the reason the Georgian Internet went dark this time around had to do with a seventy-five-year-old woman named Hayastan Shakarian, a “poor old woman” who had “no idea what the Internet is.” She had been scavenging for firewood and old copper and accidentally cut a fibre-optic cable running parallel to a railway line, severing a key Internet connection. The effect was not limited to Georgia: because of how routing was

configured in the region, Ms. Shakarian's inadvertent action also shut down the Internet in neighbouring countries. Ninety per cent of Armenia's private and business Internet users were cut off, as were many in Azerbaijan.

...

What is cyberspace? Ask most people this question and they simply shrug: for them it remains a mysterious and technological unknown that "just works." The term *cyberspace* was coined in the early 1980s by science fiction writer William Gibson, who defined it as a "consensual hallucination," and that, indeed, is how it often seems. When we log onto Twitter or Facebook through our laptops or mobile phones, we enter into what feels like an ethereal world divorced from physical reality. Our thoughts about cyberspace – if indeed these can be characterized as thoughts at all – generally begin and end with the screen in front of us. We send an email and within seconds it magically appears on a friend's BlackBerry or laptop. We text a message and it is instantly received by a colleague on the other side of the world. We start up a video on YouTube and seconds later it is streaming in high definition. We take this for granted, don't even really think about it.

But what happens in those nanoseconds as the transmission of movies or emails or Internet searches are completed? Information travels at the speed of light, and the processing power of computers is astonishingly fast. It is almost impossible to grasp that the moment a text message is sent thousands of kilometres away the information is transmitted through a complex physical infrastructure spanning multiple political jurisdictions, thousands of private companies and public entities, and numerous media of communication, from wireless radio to fibre-optic cables, like the one Hayastan Shakarian accidentally severed in Georgia.

What if it were possible to overcome the laws of space and time and follow that email, text, or tweet? What would we see? Where does the data go? Who has access to it? What happens beneath the surface of cyberspace that we don't see? Although cyberspace may seem like virtual reality, it's not. Every device we use to connect to the Internet, every cable, machine, application, and point along the fibre-optic and wireless spectrum through which data passes is a possible filter or "chokepoint," a grey area that can be monitored and that can constrain what we can communicate, that can surveil and choke off the free flow of communication and information.

...

Those constraints begin the moment we interact with the Internet, starting with the instructions that make it all work. There are millions of software programs whose instructions shape and define the realm of the possible in cyberspace, and millions more are generated every year. Software, and its codes and commands, route traffic, run programs for us, let us into the virtual worlds we inhabit. One of the unique (and disconcerting for many) features of cyberspace is that anyone can produce software that can be distributed across the Internet as a whole. Some of the most ingenious pieces of code have been written by individuals for no other reason than to get their invention "out there," to boast and take advantage of a "free" distribution network.

Not all such code is benign. Countless thousands of ever-evolving malignant programs circulate through cyberspace as viruses, trojan horses, and worms. The implications of such "malware" range from minor inconveniences to threats to privacy to debilitating attacks on national security, and some researchers believe that there is now more malware than legitimate software applications, most of it emerging too quickly for computer security

professionals to track. Malware ghettos inhabit vast and loosely connected ecosystems of insecure and outdated software programs, some of them lying dormant for years before being discovered. The progenitors prowls silently through social networking platforms, hijacking innocent people's Twitter or Facebook accounts to send phony requests to visit advertising sites or to do something more dastardly. Many of our computers may be infected by malware without our knowing it. What's worse, we pass these infections unwittingly along to friends and colleagues when we exchange information, visit malicious websites and blogs, or download documents from the Internet.

Much of the software that operates cyberspace is "closed" or proprietary, meaning that some person or company treats the code as its intellectual property. Open-source software, on the other hand, refers to code that is open to public inspection and sharing, depending on the licence. The tension between the two runs deep, and cuts across intellectual property and security issues. We may assume that closed code is relatively safe, but it is generally accepted by computer scientists that open-source code is more secure by virtue of having more "eyeballs" able to review it for potential flaws. An additional concern around closed code is the possibility that special instructions have been built in to it that might affect users without their knowledge – secret "backdoors" written into instructions by a defence or law enforcement agency, for example.

After software, the router – a device that sends information along to its destination – may be the second most fundamental choke-point in cyberspace. Most of us are familiar with the small, often frustrating, boxes with tiny antennas that give us the ability to connect to the Internet wirelessly, whether in a coffee shop, or our homes and offices. In accessing these routers, we generally choose the default security presets provided by manufacturers without

giving much thought to how easy they are to infiltrate. In a matter of minutes, armed with a \$50 Alfa AWUSO50NH USB wifi adapter (which can be purchased from Amazon) and a Linux security-testing application called BackTrack, a person without any computer engineering skills whatsoever could easily follow a set of simple instructions (laid out on YouTube, for instance) that would allow him or her to easily crack a Wireless Encryption Protocol (WEP)-enabled wifi router's password. Even simpler methods are available. Most wifi routers are shipped with default administrative passwords, accessible via a Web-based interface. Although users are cautioned to regularly change their passwords, most do not, allowing anyone to make intelligent guesses and access their routers remotely over the Internet. One website, called Router Passwords, archives known default passwords associated with router brands. How serious such vulnerabilities are was demonstrated in 2012 in Brazil, when attackers compromised 4.5 million home routers via default password hacks that changed people's DNS server settings so that when they attempted to visit websites like Google they were redirected to phony sites that looked legitimate but were in fact controlled by the attackers and contained malicious software.

Even without breaking into them, routers can leak information about us and our activities. In 2010, while mapping for its popular Street View service using its specially outfitted cars, Google collected information on wifi hotspots for use in a database it maintains to triangulate connections for mobile phones and other devices. It later emerged that Google had also collected (it claims unintentionally) payload information being secreted from unencrypted wifi routers along the way, including private information being communicated from homes and businesses. It turned out that its vehicles, outfitted with rooftop cameras and antennas, travelled up and down city streets like roving digital vacuum

cleaners sucking up telephone numbers, URLs, passwords, emails, text messages, medical records, and video and audio files sent over open wifi networks.

In 2012, Cisco provided updates to its popular Linksys EA3500 dual-band wireless router. Users were redirected away from their usual administrative interface to “Cisco Connect Cloud” instead. In doing so, however, they had to agree to new terms of service that restricted use deemed “obscene, pornographic, or offensive,” and that might “infringe another’s rights, including but not limited to any intellectual property rights.” (Cisco had also written in a clause that alluded to collecting all users’ surfing history, but removed it after considerable outrage.) These limitations on what users can do in cyberspace were put in place not by their Internet service providers or by the government, but by the private manufacturer of the hardware they used to connect to the Internet.

Also in 2012, a cyber security researcher named Mark Wuergler found that Apple’s iPhones transmitted to anyone within radio range the unique identifiers – known as MAC (media access control) addresses – of the last three accessed wifi routers. He cross-checked that information against a publicly accessible database of MAC addresses to pinpoint their locations on a map. Wuergler then created an application called Stalker to make it easier to harvest and analyze unintentionally leaked information – passwords, images, emails, and any other data transmitted by mobile phones and wifi routers. The information collected by Stalker contained the names of specific businesses regularly frequented, or friends and colleagues who are regular chat buddies. That information could be used to deceive someone into revealing further data which, in turn, could be used to undertake electronically based attacks.

The Citizen Lab uses a similar network analyzing tool, Wireshark, to sniff out hidden details of Internet traffic, though we do so only with the permission of those we monitor. Wireshark data has

allowed us to see questionable connections being made to remote servers and evidence of malicious activity, as we found during our GhostNet probe. We have used the tool in workshops to demonstrate how much information can be gathered remotely without an Internet user's knowledge. Using Wireshark and connecting to a wifi network in a hotel, for example, one can collect information on who is attending a private meeting in a room down the hall (based on computer name data sent over the Internet) and sometimes usernames and passwords (if they are sent unencrypted). It is possible to collect data on all of the sites being visited and data downloaded by users in the room, the content of private chats, and updates to Twitter and Facebook accounts (again, if the user's communications are not encrypted).

We also use a tool called Nmap to scan networks and map the computers connected to them, which ports are open on those computers, what operating systems are used, et cetera. With Wireshark and Nmap employed together, we can precisely map the computers and devices logged onto a network (including all known vulnerabilities on those computers and devices), and collect much of what is being communicated by the people using those computers. All of this information can be collected – without the users ever noticing – by someone connecting to the same network a few metres down the hall using a few freely available open-source tools. Examples like these show how the multiplying access points into cyberspace can create unintentional vulnerabilities that may expose us to security and privacy risks.

• • •

We take cyberspace for granted. We assume that its basic modus operandi – uninterrupted connectivity to a shared communications environment – is always stable. That assumption is wrong.

Cyberspace is a highly dynamic ecosystem whose underlying contours are in constant flux. One of the most important recent changes has come about with the gradual movement away from searching the World Wide Web to a “push” environment where information is delivered to us instead, mostly through applications and services. A major impetus behind this shift has been the popularity of mobile devices, especially the Apple iPhone. Web browsers are functionally constrained by the smaller screens and other limitations of mobile devices, which has led to the popularity of applications that deliver specially tailored information to users instead. So, whereas in the past we might have visited the *New York Times* website via our browser, today a growing number of us download the *Times* app instead, signing off on another terms of service licence agreement in the process, and sharing with yet another third party a potentially far greater amount of personal data connected to our mobile phone. Of course, what can be “pushed out” can also be “pulled back” by companies, or turned off at the request of governments. Apple’s iPhone, for instance, has a built-in remote “wipe” functionality that can permanently disable or erase the device and all of its apps.

When we communicate through cyberspace, our data is entrusted to the companies that own and operate the hardware, the applications and services, and the broad infrastructure through which our communications are transmitted and stored. These companies are the intermediaries of our Internet experiences, and what they do with our data can matter for how we experience cyberspace, and what we are permitted to do through it. They are critical agents in determining the rules of the road by virtue of the standards they insist upon, the operating decisions they take, and the constraints they impose on users. This is especially important as the volume of data they control becomes ever greater, ever more potentially lucrative in the global information economy.

The end-user licence agreements, terms of service, and other warranties we sign with these companies define what they can do with our data. Unfortunately, few users bother to read, let alone understand, them. It is hard not to be sympathetic. Unless one has an advanced legal degree, these documents are intimidating: tens of thousands of words in fine print, with exceptions and caveats that provide enormously wide latitude for what companies can do. Faced with this word-soup, most of us just click “I agree.” What we are agreeing to might surprise us. Skype users, for instance, might be alarmed to find out that when they click on “I agree” to the terms of service they are assigning to Skype the right to change these terms at any time, at Skype’s discretion, and without notice. Skype does not inform users about whether and under what conditions it will share user data with law enforcement or other government agencies. Users might not know that while they can stop using Skype, they cannot delete their accounts: Skype does not allow it.

The Internet is sometimes described as a massively decentralized and distributed “network of networks,” a virtual place where information from everywhere is concentrated and accessible to all, an egalitarian thing of beauty. From one perspective, this description accurately characterizes its architecture. But within this network of networks there are critical chokepoints: a tangible, physical infrastructure that includes the hardware, software, cables, even the electromagnetic spectrum that exists in definable, real space. There are also regulatory and legal chokepoints: the ways in which cyberspace is structured by laws, rules, and standards that can facilitate forms of control. Mobile forms of connectivity, now the central method of communicating in cyberspace, are a case in point. The mobile industry is controlled by manufacturers of “closed” devices, handsets whose owners are prohibited from opening (“jailbreaking”) or fiddling with their insides at the risk of warranty violations. Closed or proprietary software (with the exception of

Google's open-source Android operating system) means that the millions of lines of instructions that run a mobile phone's system are restricted to everyone but the company that sells the software. They operate through networks owned and serviced by a small number of ISPs and telecommunications companies (sometimes only one or two, depending on the region or country), and they function according to government-issued licences that set the conditions under which people can use the wireless spectrum. What from one perspective (that of the average user) looks like an ephemeral network of networks, from another (that of people in positions of authority) looks more like a tangible system of concrete controls through which power can be exercised and the nature of communications shaped for specific political or economic ends.

• • •

It is important to understand the political architecture of cyberspace, because the companies that own and operate its infrastructure, applications, and devices are under increasing pressure from a variety of quarters to police the networks they manage: from the technical demands of managing increasingly complex types of communication flows like bandwidth-sucking video streams; from lucrative market opportunities to repack and sell user data; from regulations passed down by governments to corporations to manage content and users. The latter are especially noteworthy because the Internet crosses political boundaries, and many companies have operations in multiple national jurisdictions, some of which do not respect the rule of law or basic human rights and whose policing of the Internet lacks transparency. To operate in some jurisdictions search engines, mobile carriers, and other Internet services are required to filter access to content deemed objectionable by host governments, turn services on and

off in response to crises, push intimidating mass messages onto citizens living in certain regions, cities, or territories, and/or share information about users with state security services. More often than not the companies comply.

The cyberspace experience can vary dramatically depending on what application or device we use, which Internet café or hotspot we log on from, which ISP we contract with, and, most fundamentally, which political jurisdiction we connect from. Without the aid of special anti-censorship software, an Internet user in China is unable to connect to Twitter or Facebook, while a user in Pakistan cannot view YouTube. Users in Thailand cannot access videos on YouTube deemed insulting to the royal family. A user of the ISP du Telecom in the United Arab Emirates cannot access information about gay and lesbian lifestyles. (Using filtering technology produced by the Canadian company Netsweeper, such content is censored by du.) Indonesian users of BlackBerry devices are not able to access thousands of websites deemed pornographic and blocked by Research in Motion (RIM). Individuals living in volatile Kashmir are not able to access Facebook. According to ONI (OpenNet Initiative – a collaborative partnership of the Citizen Lab; the Berkman Center for Internet & Society at Harvard University; and the SecDev Group in Ottawa), dozens of governments now insist that ISPs operating in their political jurisdictions implement Internet censorship and surveillance on their behalf.

Internet filters and chokepoints can have bizarre collateral impacts on users' Internet experiences around "upstream filtering," cases where data transit agreements, or "peering," made between ISPs in separate countries can have spillover effects on Internet users in each others' countries. In 2012, ONI discovered that users in Oman were not able to access a large number of websites with Indian-related content (mostly Bollywood movies and Indian music). The source of the censorship, however, was not in Oman itself nor was it

demanded by the government (for whom the sites in question were not controversial). Rather, it was the Indian ISP Bharti Airtel, with whom the Omani ISP, Omantel, has a peering arrangement.

This kind of collateral impact of Internet controls has a long history. In 2005, ONI found that when the Canadian ISP Telus blocked subscriber access to a website set up by a labour union intending to publicize its views about a dispute with Telus, it also unintentionally blocked access to over 750 unrelated websites. In 2008, the Pakistan Ministry of Information ordered Pakistan Telecom to block access to YouTube because of films uploaded to the site that purportedly insulted the Prophet Muhammad. In carrying out this order, Pakistan Telecom mistakenly communicated these routing instructions to the entire Internet, shutting down YouTube for most of the world for nearly two hours.

• • •

Most of the filtering described above takes place at the level of ISPs, the companies users contract with to get their basic connectivity. But there is a deeper layer of control, one that stretches down into the bowels of cyberspace: Internet Exchange Points (IXPs). While most users are familiar with ISPs, few have ever heard of IXPs. There are several hundred IXPs around the world: usually heavily guarded facilities with the level of security one encounters at an airport or defence installation. If you've ever wondered how it is that your email reaches your friend's email account with a completely different company, IXPs are the answer. It is here that traffic is passed between the networks of different companies – through border gateway protocols (BGP) exchanged between ISPs – and IXPs are the key strategic locations for the interception, monitoring, and control of large swathes of Internet communications. (In the early 2000s, I toured an IXP in downtown Toronto

and saw row upon row of high-tech equipment, endless servers stacked on several floors. Down one long hallway there were hundreds of what appeared to be randomly distributed red tags attached to the equipment. I asked the tour guide, "What are the red tags?" He replied nonchalantly, "Oh, those are the wiretaps," and moved on.)

In 2002, Mark Klein, a twenty-year veteran technician with AT&T, was working at an IXP in San Francisco. He became suspicious after noticing some unusual activity in a "secure room" marked 641A. Klein was working in an adjacent area and had been instructed to connect fibre-optic cables to cables exiting from the secure room. He was not allowed to enter the room, and the people there were not the type of workers with whom Klein enjoyed lunch and coffee breaks. They kept to themselves and seemed to have special privileges. Later, Klein learned from his colleagues that similar operations were observed by engineers at other AT&T facilities across the United States.

Klein's suspicions eventually led to a class action lawsuit by the Electronic Frontier Foundation (EFF) against AT&T, alleging that the company had colluded with the National Security Agency (NSA) outside of the rule of law. As it turned out, inside room 641A was a data-mining operation involving a piece of equipment called Narus STA 6400, known to be used by the NSA to sift through large streams of data. The choice of location was significant. Because of the complex routing arrangements that govern the flow of traffic through cyberspace, many smaller ISPs sublease their traffic through AT&T – a globe-spanning "Tier 1" telecommunications company – and a large proportion of global communications traffic flows through its pipes. The AT&T-operated IXP in San Francisco is one of the world's most important chokepoints for Internet communications.

The IXP is a chokepoint for not only international traffic; it handles a large volume of domestic U.S. communications as well.

The NSA is prohibited from collecting communications from American citizens, and the data-mining operation at the AT&T facility strongly suggested that prohibition was being ignored. The EFF class action lawsuit took AT&T and another IXP operator, Verizon, to task for their complicity with what turned out to be a presidential directive instructing the NSA to install the equipment at key IXPs in order to monitor the communications of American citizens. In 2008, as the lawsuit dragged on, the Bush administration took pre-emptive action by introducing a controversial amendment to the Foreign Intelligence Services Act (FISA), giving telecommunications companies retroactive immunity from prosecution if the attorney general certified that surveillance did not occur, was legal, or was authorized by the president. This certification was filed in September of 2008 and shortly thereafter, the EFF's case was dismissed by a federal judge citing the immunity amendment. (Presidential candidate Barack Obama surprised many of his supporters by backing the FISA Amendment Act, and his administration has vigorously blocked court challenges against it ever since.) Although the full scope of the NSA's warrantless wiretapping program (code-named "Stellar Wind") is classified, William Binney, a former NSA employee who left the agency in protest, estimates that up to 1.5 billion phone calls, as well as voluminous flows of email and other electronic data, are processed every day by the eavesdropping system stumbled upon by Klein.

IXmaps, a research project at the University of Toronto, raises awareness about the surveillance risks of IXPs, particularly for Canadians. The project uses trace-routing technology to determine the routes discrete bits of information (or "packets") take to reach their destination over the Internet. In one example, IXmaps detailed the route of an email destined for the Hockey Hall of Fame in downtown Toronto and originating at the University of Toronto a few miles away. The email crossed into the United

States, was peered at an IXP in Chicago, and was probably exposed to one of the NSA's warrantless surveillance systems rumoured to be located at the facility. Known as boomerang traffic, this type of cross-border routing is a function of the fact that there are eighty-five IXPs in the U.S., but only five in Canada. Routing arrangements made by Canadian ISPs and telecommunications companies will routinely pass traffic into the U.S. and back into Canada to save on peering costs, subjecting otherwise internal Canadian communications to extraterritorial monitoring.

• • •

One of the long-standing myths about cyberspace is that it is highly resilient to disruption. For those of us who have laboured over Internet downtimes, email failures, or laptop crashes, this may seem like a fanciful idea. But the resiliency of cyberspace does have some basis in the original design principles of the Internet, whose architecture was constructed to route information along the most efficient available path and to avoid disruption in the event of a natural disaster (or nuclear attack). This resiliency was demonstrated in the aftermath of Hurricane Sandy in October 2012, which devastated the U.S. eastern seaboard and caused mass power outages, including the loss of local Internet and cell-phone connectivity. The network-monitoring company Renesys showed that the storm had collateral impacts on traffic as far away as Chile, Sweden, and India – but mostly in a positive sense: traffic destined for New York City that would have failed as a consequence of the storm was manually rerouted along alternative paths by savvy network engineers.

However, there are also many characteristics of cyberspace that demonstrate fragility and a lack of resiliency; Hayastan Shakarian's mistaken severing of an underground cable in Georgia to name

one. It may come as a surprise that the same type of cables that Shakarian accidentally unearthed traverse the world's lakes and oceans, and bind cyberspace together in a very material sense. Undersea cables are one of the links that connect today's cyberspace to the late Industrial Revolution. The first such cables were laid in the late nineteenth century to facilitate telegraph traffic over long distances. Early designs were prone to failure and barely allowed the clicks of a telegraph exchange to be discerned across small bodies of water like the English Channel, but over time innovations in electronics and protective cable sheathings allowed the undersea cable industry to flourish. (This growth led to a dramatic increase in international telephone calls, and a new market for the sap of gutta-percha trees, which was used to coat and protect the cables until the mid-twentieth century.) Although international telecommunications have been supplemented with microwave and satellite transmissions, a surprisingly large volume of data still traverses the world through cables crossing the Atlantic and Pacific oceans, and major bodies of water like the Mediterranean Sea.

Due to the staggering costs involved, companies often share the same undersea cable trenches and sometimes competing companies even share the same protective sheathing. This makes those trenches highly vulnerable to major disruption. In a May 2012 article published on the website Gizmodo, provocatively titled "How to Destroy the Internet," the author details the physical elements of the Internet that could be easily targeted. He provides a link to a document alphabetically listing every single cable in the world, and its landing stations. While there are hundreds of cables, the total is not astronomical – and probably a lot fewer than what most people might expect for a network as vast as the global Internet. Among them is ACS Alaska-Oregon Network (AKORN), with its landing points in Anchorage, Homer, and Nikiski, Alaska,

and Florence, Oregon; the Gulf Bridge International Cable System, with its landing points in Qatar, Iraq, Bahrain, Saudi Arabia, Oman, Iran, the United Arab Emirates, Kuwait, and India; and at the end of the long list, Yellow/Atlantic Crossing-2 (AC-2), which connects New York City to Bude in Cornwall, U.K. The author goes on to explain how many of the cables' onshore landing stations are sometimes "lying out on the sand like an abandoned boogie board," and how the cables could be severed with a few swings of an axe. Severing cables in this way at landing stations in only a few select locations – Singapore, Egypt, Tokyo, Hong Kong, South Florida, Marseilles, Mumbai, and others – could wreak havoc on most of the world's Internet traffic.

The 2006 Hengchun earthquake, off the coast of Taiwan, affected Internet access throughout Asia, and in 2008 two major cable systems were severed in the Mediterranean Sea. The cause of the severed cables is unknown, but some experts speculated that the dragging of a ship's anchor did the job. But a review of video surveillance taken of the harbour during the outage period showed no ship traffic in the area of the severed cable. Others suggested it could have been a minor earthquake, causing a shift in the ocean floor, but seismic data didn't support this conjecture. Whatever the cause, such cuts to cables are fairly routine: Even in their trenches, undersea cables are pushed to and fro by currents and constantly rub against a rough seafloor. In the case of the 2008 Mediterranean incident, the damage was severe: there were disruptions to 70 percent of Internet traffic in Egypt and 60 percent in India, and outages in Afghanistan, Bahrain, Bangladesh, Kuwait, the Maldives, Pakistan, Qatar, Saudi Arabia, and the United Arab Emirates. Nearly 2 million users were left without Internet access in the U.A.E. alone. Connections were not restored until a French submarine located the severed cables and brought them to the surface for repair.

Prior to the introduction of fibre optics, undersea cables were occasionally wiretapped by attaching instruments that collect radio frequency emitted outside the cables. During the Cold War, both the United States and the Soviet Union built special-purpose submarines that would descend on cables deep in the ocean and attach inductive coils to collect emissions. In his book *Body of Secrets*, historian James Bamford describes in detail Operation Ivy Bells in the early 1970s, in which the NSA deployed submarines in the Sea of Okhotsk to tap a cable connecting the Soviet Pacific Naval Fleet base in Petropavlovsk to its headquarters in Vladivostok. Specially trained divers from the USS *Halibut* left the submarine in frigid waters at a depth of 120 metres and wrapped tapping coil around the undersea cables at signal repeater points, where the emissions would be strongest. Tapes containing the recordings were delivered to NSA headquarters, and were found by analysts to contain extraordinarily valuable information on the Soviet Pacific Fleet. Several other submarines were later built for such missions, and deployed around the Soviet Union's littoral coastline and next to important military bases. When fibre-optic technology (which does not emit radio frequencies outside of the cable) was gradually introduced, the utility of such risky operations diminished. However, some intelligence observers speculate that U.S. and other signals intelligence agencies have capabilities to tap undersea fibre-optic cables by cutting into them and collecting information through specifically designed splitters.

• • •

Like undersea cables, satellites illustrate the fragile nature of cyberspace. In 2009, a defunct and wayward Russian satellite collided with an Iridium low Earth orbit satellite at a speed of over 40,000 kilometres per hour. The collision caused a massive cloud of

space debris that still presents a major hazard. NASA's Earth observation unit tracks as many as 8,000 space debris objects of ten centimetres or more that pose risks to operational satellites. (There are many smaller objects that present a hazard as well.) The Kessler Syndrome, put forward by NASA scientist Donald Kessler in 1976, theorizes that there will come a time when such debris clouds will make near-Earth orbital space unusable. Although undersea fibre-optic cables provide the bulk of transit for global communications, they cannot sustain the entire load. A scenario such as the Kessler Syndrome, were it to come true, would end global cyberspace as we know it. Scientists have very few realistic solutions for cleaning up space debris.

Space is also an arena within which state intelligence agencies exercise power over the Internet. Although the Apollo missions were publicly justified on the basis of advancing human curiosity and science, the first missions into space actually had specific military and intelligence purposes. Since the 1960s, the superpowers have been developing globe-spanning satellites that are used for optical, infrared, thermal, and radar reconnaissance purposes. The Americans built a fleet of specially designed satellites whose purpose is to collect signals intelligence (sigint). Some sigint satellites operate in geostationary orbit 36,000 kilometres from the Earth's surface, and are used to zero in on radio frequencies of everything from microwave telephone signals to pagers and walkie-talkies. Such geostationary sigint satellites deploy huge parabolic antennas that are unfolded in space once the satellite is in position, with the signals being sent to NSA listening stations located in allied countries like Australia (Pine Gap), and Germany (Bad Aibling). Because the satellites operate in deep space, and radio signals travel in a straight line, radio frequencies can be collected efficiently and with little degradation. (Other sigint satellites take unusual orbits and can reportedly hover over regions of interest for longer periods and at lower altitudes.)

The NSA also operates sigint collection facilities at ground stations whose mission is to collect transmissions from civilian communications satellites. Typically, these enormous interception terminals, which look like giant angled birdbaths, are located in secure areas proximate enough to terrestrial transmission points to function properly. For example, one of the key signals intelligence stations in Canada is at the Canadian Forces Station Leirtrim, just south of Ottawa, strategically positioned to intercept diplomatic communications moving in and out of the nation's capital.

Signals intelligence gathering is highly secretive, but it is a world we should all get to know better. Originally, the objects of sigint operations were other states' military and intelligence agencies: ballistic missile-test telemetry or operational instructions sent by high-ranking Politburo members. As the Cold War came to a close, however, this bipolar conflict atomized into a multitude of national security threats, some of which emanate from transnational terrorist groups and organized crime, and the scope of sigint operations became much broader and more widely dispersed across global civil society. As the volume of data flowing through global networks is exploding in all directions, and the tools to undertake signals intelligence have become more refined, cheaper, and easier to use, the application to cyberspace is obvious.

• • •

Although cyberspace is often experienced as an ethereal world separate from physical reality, it is supported by a very real infrastructure, a tangible network of code, applications, wires, and radio waves. Behind every tweet, chat message, or Facebook update, there is also a complex labyrinth of machinery, cables and pipes buried in trenches deep beneath the ocean, and thousands of orbiting satellites, some the size of school buses. In addition to

being complex and fragile, this physical infrastructure contains a growing number of filters and chokepoints. Pulling back its layers is like pulling back curtains into dark hallways and hidden recesses, which, it turns out, are also objects of intense political contests.

There is another component of cyberspace, separate from its physical infrastructure, but that is also growing in leaps and bounds and becoming a critical part of the domain: the data. Information related to each and every one of us (and everything we do) is taking on a life of its own. It, too, has become an object of geopolitical struggle. Every call we make, every text and email we send, increasingly everything we do as we go about our daily lives, is recorded as a data point, a piece of information in the ever-expanding world of “Big Data” that is insinuating itself deeper and deeper into our lives and the communications environment in which we live.