23

The geopolitics of internet control

Censorship, sovereignty, and cyberspace

Ronald J. Deibert

What is the impact of the internet on state sovereignty, and in particular on states' ability to control information flows across their borders? Whereas once the internet was presumed to be a borderless world of free-flowing information, today countries and corporations alike are carving it up in a bewildering array of filtered segments, often with major unintended consequences. The motivations for these practices range widely, from concerns over national security, cultural sensitivities, and protection of social values, to rent seeking and the protection of economic monopolies. Whereas once it was conventional wisdom to believe that the internet's technological infrastructure was immune to control, today states and corporations are applying an ever-increasing level of skill and technological sophistication to precisely that mission. The result is that rather than being a single seamless environment, the internet a user connects to and experiences in Canada is far different than an internet a user experiences in Iran, China, or Belarus. This chapter provides an overview of the geopolitics of internet control, and in particular state efforts to control information flows across borders, with comparative data from over 22 countries.

In early 2007, the online mapping service Google Earth provided a feature on the ongoing political crisis in the Darfur region of Sudan. Not long afterwards, however, an aid worker based inside Sudan reported not being able to properly load the map, receiving an error message in his browser stating "This product is not available in your country." Upon further inspection, the source of the inaccessibility was Google itself—filtering access to its own services based on the "geolocation" of the computer's IP address making the request. Google was not permitting IP addresses based within Sudan from connecting to its service in order to comply with U.S. export restrictions against the sale or export of informational products to the country (Geens, 2007).

Earlier the same year, Tunisian authorities filtered the popular video-streaming service, DailyMotion. DailyMotion is known to carry a wide range of political videos, including many satirical videos of the Tunisian government's record on human rights. Many inferred that Tunisia had blocked the website because of those videos, following its known track record of blocking access to opposition and human rights websites (Reporter Without Borders, 2007). However, Tunisia uses (but does not openly admit to doing so) the U.S. commercial filtering product, Smartfilter, to block its citizens' access to information (OpenNet Initiative, 2005a). DailyMotion was, perhaps mistakenly, categorized within the Smartfilter database as "pornography"—a category apparently

selected by Tunisia for blocking. After reports of the DailyMotion block surfaced, Smartfilter apparently corrected the categorization error, and access to the DailyMotion website from within Tunisia was gradually restored.

The source for much of the evidence and illustrations used in this chapter comes from the research of the OpenNet Initiative (ONI)—collaboration among the Citizen Lab at the University of Toronto, the Berkman Centre for Internet and Society at Harvard Law School, the Cambridge Security Programme, U.K., the Oxford Internet Institute, and partner non-governmental organizations (NGOs) worldwide.[1] The aim of the ONI is to document empirically patterns of internet censorship and surveillance worldwide using sophisticated means of technically interrogating the internet directly. The ONI's tests are carried out both remotely from North America and the U.K., and in-field by dozens of local researchers. Our reports over the last several years have documented a disturbing increase in the scale, scope, and sophistication of internet censorship practices worldwide.[2] This chapter summarizes some of the main findings of this research and draws connections to wider implications for global politics, security, and human rights. The main questions addressed by this chapter are: how many states are filtering access to information on the internet? What are the types of content that these states are targeting for filtering? What are the most effective methods used by states that filter? What is the range of transparency and accountability practices among states that filter? Are states open about their practices? And, what are some of the wider implications of these practices? As will be described in this chapter, the picture of the internet that emerges from this research is of a hotly contested and deeply politicized realm.

Beneath the surface of internet communications

What happens to a request when a user clicks on a link to a website or sends an e-mail? For most surfers, the internet experience begins and ends with what happens on the computer screen in front of them. However, if surfers follow that e-mail or web request as it leaves a computer and passes down the fiber optic cable to the servers and routers of a local internet service provider (ISP), through the internet exchange points (IXPs), international gateways, and on to the undersea trunk cables of tier 1 telecommunication companies, they will find a complex and largely hidden infrastructure of filters and chokepoints.

Most people assume that the internet's vast infrastructure is an open, decentralized, network of networks through which information flows freely along a shared routing protocol. While this description has some basis in the historical evolution of the internet, and captures parts of what makes it unique, it also obscures some of the details that structure internet communications beneath the surface. While it is true that there is no single node through which all traffic passes on the internet, and thus no form of centralized control, there are thousands of nodes that parse out and filter information and act as gateways. Each of these nodes and gateways—from routers to IXPs to autonomous systems—present opportunities for authorities to impose order on internet traffic through some mechanism of filtering and surveillance. Some of this control takes place for technological reasons; some of it takes place for cultural, political, and economic reasons. Instead of a network of networks, therefore, it is perhaps more accurate to characterize the internet as a network of *filters and chokepoints*.

The means by which content is blocked or filtered on the internet vary

widely in terms of complexity, effectiveness, and intent. Furthermore, not all of the means by which states attempt to control the internet are technological. In some cases, regulations are employed to supplement technical controls, which can create a climate of self-censorship among internet users. The following section defines some of the central terms associated with internet content filtering and surveillance before turning to specific examples of accountability and transparency issues.

Internet content filtering is a term that refers to the techniques by which control is imposed on access to information on the internet (Deibert and Villeneuve, 2004). Content filtering can be divided into two separate techniques: address blocking techniques and content analysis techniques. Address blocking techniques refer to particular router configurations used to deny access to particular internet protocol (IP) addresses and/or domain names, or specific services that run on particular port numbers. For example, a state may run a blocking filter at the international gateway level that restricts access from within the country to websites that are deemed illegal, such as pornographic or human rights websites.

Content Analysis refers to techniques used to control access to information based on its content, such as the inclusion of specific keywords on a website or the address of a URL. Because parsing mechanisms employ keywords to block access, they are often the source of mistaken or unintended blockages. Unintended blocking can occur as a result of IP based blocking as well, however, as it is not uncommon for many domain names to share the same IP address. Filtering that aims to block access to a specific website by blocking its IP address, in other words, can result in the collateral filtering of potentially thousands of unrelated sites sharing the same IP.

Depending on need and circumstance, different approaches to filtering can be implemented:

- Inclusion filtering: users are allowed to access a short list of approved sites, known as a "white list," only. All other content is blocked.
- Exclusion filtering: restricts user access by blocking sites listed on a "black list." All other content is allowed.
- Content analysis: restricts user access by dynamically analyzing the content of a site and blocking sites that contain forbidden keywords, graphics, or other specified criteria.

The mechanisms used to do these types of filtering vary considerably. Routers act as junctions between networks, passing information packets back and forth, and thus routers are the main (though not only) nodes where such blocking takes place in the form of instructions written into the routing tables. However, filtering software can be implemented into virtually any node throughout the internet's system. As a consequence, the level at which filtering can be implemented varies widely too. Filtering can take place on an individual's personal computer, an office local area network (LAN), an internet café, an ISP, a wireless network, an SMS system, at the backbone or international gateway level, or some combination of all of these levels. Not surprisingly, national-level internet content filtering can vary dynamically, and across ISPs within a single country (Anderson and Murdoch, 2007).

Although filtering traditionally takes place by blocking requests for information from either reaching their destination or returning the requested information at information chokepoints, other non-filtering mechanisms can be employed that achieve the same ends. After all, filtering is simply denial of access to information.

As is described below, new forms of blocking access to information are emerging based on the use of distributed denial of service attacks. Such attacks bring web servers down by overwhelming them with requests for information, thus "filtering" information at its source and denying access to all users equally. The same type of denial of service can (and occasionally does) take place by cutting off power to the building where web servers are located, or misconfiguring routing tables to cause what appear to be network errors, but which in fact are deliberate attempts to shut off communications at the source.

As the Google Earth example demonstrates, filtering can also take place through reverse geolocation—that is, the server hosting websites can refuse to take requests from users based on the geographical origin of their computer's IP address. The ONI has documented numerous instances of this type of reverse geolocation filtering, including by the website georgewbush.com during the 2004 U.S. Presidential Elections (ONI, 2004).

Methods of investigating censorship

Although filtering practices are widespread, knowledge of their use by states has tended to be limited. In part, this is a function of a lack of accountability and transparency among states that block access to information. In part, however, it is also a function of the lack of empirical evidence about such practices. Up until recently, the majority of reports on internet filtering tended to emerge from users, news reports, or advocacy organizations. Not surprisingly, they tended to be unsystematic and sometimes even unreliable. Moreover, because of the complex and varied ways in which filtering can be implemented, as noted earlier, reports have often been made in error or have contained contradictory information.

The aim of the ONI has been to overcome these shortcomings by developing a systematic way to investigate empirically internet filtering practices from within state borders over an extended period of time. The project employs a unique methodology that combines in-field investigations by partners and associates who travel to or live in the countries concerned, and a suite of technical interrogation tools that probe the internet directly for forensic evidence of content filtering and filtering technologies.[3] These tools work from the "inside out" of the internet, probing parts of the information infrastructure not generally apparent to the average user. The methods range from automating connecting requests to servers hosting websites simultaneously from within the country under investigation and a control location in a non-filtered location, to using tracing and other network mapping tools to interrogate the location of and technologies used to do the filtering. Tests for accessibility to internet content were based on categorized lists of websites.[4] These categories were meant to cover as comprehensively as possible the likely targets for filtering by states while allowing for as precise as possible identification of content categories singled out for filtering. While most states that filter target pornographic content, as will be shown later a wide range of non-pornographic, political content—such as opposition parties or minority rights, for example—is now being targeted as well by several states.

This method allows for a comprehensive picture of internet content filtering in a particular country by probing all aspects of the national information infrastructure (internet cafés, ISPs, wireless networks, backbone gateways) and over an extended period of time testing accessibility in both English and local languages to lists of

thousands of websites in each of these categories.[5]

Since 2002, the project has produced detailed reports on 11 countries—Belarus, Yemen, Tunisia, Burma, Singapore, Iran, China, Bahrain, United Arab Emirates, Vietnam, and Saudi Arabia. More recently, in 2006 the ONI conducted extensive tests over several months in more than forty countries worldwide. The following sections highlight some of the main trends and findings emerging from this research.

The globalization of online censorship

In 2002, only a handful of countries were known to engage in internet content filtering, most prominently China, Iran, and Saudi Arabia. By 2007, 26 of 40 examined countries were found to engage in internet filtering practices to some degree.

China is still the world's most notorious and sophisticated censoring regime (ONI, 2004, 2005a, b, c, d; Dowell, 2006; Li, 2003; Li, 2004). Its filtering system comprises multiple levels of legal regulation and technical control, the latter implemented primarily at the backbone level using specially configured Cisco routers. The system involves numerous state agencies and thousands of public and private personnel, and a dense web of ever-thickening legal restrictions.

The range of information that China seeks to limit and control from within its borders is broad. China targets content for filtering across every major category tested, including human rights, opposition and independence and secessionist movements, minority faiths, pro-democracy groups, search engines, free e-mail and webhosting services, anonymizers and circumventors, pornography and sexually explicit material, and others.

However, China is not alone. Although many countries justify their censorship practices as a way to block access to pornography or other culturally sensitive material, our research has documented a large and growing swathe of content beyond pornography that is targeted for filtering. At least 14 countries blocked access to content that spans the major categories of *political*, *social*, and *conflict/security* content, including Burma, China, Ethiopia, Iran, Oman, Syria, Thailand, Tunisia, United Arab Emirates (UAE), Uzbekistan, Vietnam, Pakistan, Saudi Arabia, Sudan, and Yemen (See Figure 23.1).

Some of the countries in which we found evidence of content filtering in each of these major categories began by blocking only a few select sites in one category, usually pornography. After a period of time, however, the scope of content targeted for filtering began to increase to other content areas. In Thailand, for example, what started out as an effort to block pornography has been gradually broadened to include politically sensitive websites as well, particularly since the September 2006 military coup. In addition to pornographic content, Thailand blocks access to the popular video streaming service, YouTube.com, ostensibly in response to a single video posted on the service satirizing the deposed King. Pakistan began filtering websites that contain imagery offensive to Islam, and now targets all sites related to the Balochistan independence movement as well. The Thai and Pakistan cases are illustrative of what may be a more general trend: that is, once the tools of censorship are put in place, the temptation for authorities to employ them secretly for a wide range of ulterior purposes may be large—particularly in circumstances where there is little civilian oversight or accountability—a phenomenon we refer to as internet censorship "mission creep."

A number of other countries were found to be engaged in less pervasive forms of internet filtering, typically concentrated

around a single content area or contentious internet service. For example, in addition to blocking some gambling and pornographic sites, ISPs in South Korea block access to all websites related to North Korea. India blocks access to websites related to extremist and militant groups, particularly those associated with Hindu and Islamic extremism. A number of Middle Eastern and Gulf Countries, including Syria, Jordan, UAE, Bahrain, and Saudi Arabia, block access to the entire Israeli (.il) domain (see also Warf and Vincent, 2007). Though having strict controls over traditional media and heavy penalties for libel, Singapore blocks access only to a small handful of pornographic websites (see also Rodan, 1998). Following the Thai and Pakistani examples above, we might hypothesize that over time these states will likely use their filtering systems to block a growing body of content.

Increasing censorship sophistication

Not surprisingly, the methods used to do internet content filtering have become more sophisticated, as states and the firms that sell censorship and surveillance technologies continually refine them. There are several examples of increasing sophistication. First, authorities are becoming increasingly adept at targeting newly developed modes of communication, such as blogs, SMS, chat, and instant messaging

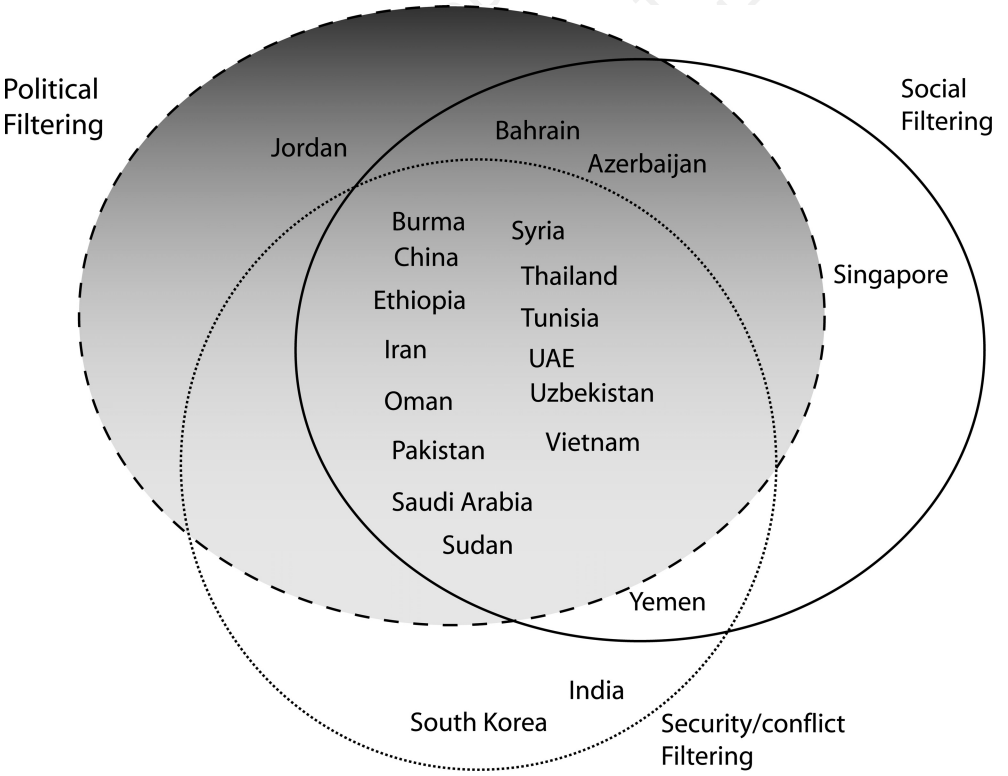

Figure 23.1 Content filtering by major category.
Source: Faris and Villeneuve, 2006.

protocols, and voice over internet protocol (VOIP) services. In the past, such newly devised methods of information sharing could be used as a means to circumvent internet censorship. However, today authorities are becoming more adept at targeting new media and developing methods particular to such services. Second, although content filtering is prone to overblocking and error, there are examples where authorities have been able to use such technologies with precision.

A good example is China's targeting of the specific string of codes embedded in the URL of the Google cache function. The latter is a service provided by Google whereby users can connect to archived information from websites stored on Google's servers, rather than on the servers of the original website. The service was designed to provide a way to access information through redundancy, but it is also a very simple and effective way to get around content filtering. Since users connect to Google servers rather than to the blacklisted servers, they bypass the content filters. Upon learning of this technique, China implemented a blocked string on their backbone/gateway routers that prevented *any* use of the Google cache function from within China.

A third example of increasing sophistication of content filtering is the targeting of local languages and websites of opposition movements and dissidents particular to a specific national context. Tests from within China comparing the top 100 Google search results for keywords in English and Chinese show a very significant disproportionate amount of keywords are filtered when they are searched for in Chinese as opposed to English (ONI, 2005b). For example, a search for the terms "Chinese Labor Party" in Chinese yields a 93 percent inaccessible rate when compared to the same search performed in English, which yields only a 20 percent inaccessible rate. Iran, in 2005, showed a similar relationship among English and local language filtering (ONI, 2005c). In the case of Iran, many of the blocked websites in various categories had a higher percentage of inaccessibility in Farsi as opposed to English. Overall, 80 percent of the Farsi-language websites tested were inaccessible whereas 45 percent of English-language sites were inaccessible. Such *localization* filtering—where "international" sources of information are left accessible while local variants are blocked—may at first seem counterintuitive. However, there are two potential explanations. First, localization filtering targets those groups that matter most to regime stability and power, such as local opposition movements and dissident groups presenting contentious information in languages spoken by citizens within the country. Second, the disproportionately open access to English-language international sites can give the impression that access to global information is wide open, particularly to foreign journalists who do not speak local languages. Authorities can point to contentious human rights and news sites and say that they allow access to information while blocking relatively more obscure sites from a global perspective that matter most in local politics.

The tests conducted across 40 countries in 2006 provided further confirmation that state content filtering tends to concentrate on local content and websites. Table 23.1 shows the percentage of websites blocked by country in the local and global content categories respectively. For each country, two baskets of websites were used for comparison: a local list, which includes categorized websites related to the particular context of each country in question; and a global list, which is a control list of categorized websites tested for accessibility in every country. The local list contains mostly local-language content of each country in

question (e.g., Farsi for Iran) while the global list contains English-language content. The percentage of blocked websites in the local category was higher than in the global category for many countries in which filtering was found. When pornographic-related content is removed (which tends to be mostly global in character and filtered as a default by many countries using commercial filtering applications), the percentage of local content targeted for filtering is even higher. Among countries found to be engaging in content filtering, UAE, Bahrain, China, Ethiopia, India, Iran, Korea, Libya, Myanmar, Pakistan, Saudi Arabia, Syria, Thailand, Tunisia, Uzbekistan, and Vietnam all blocked relatively more local- than global-related content. What this suggests is that states who filter the internet tend to concentrate on social and political content that matters most within their local country context.

Use of commercial filtering technologies

The increased sophistication of internet content-filtering practices can be attributed, in part, to the services provided by western (mostly U.S.-based) software and internet service firms. Whereas once the best and brightest of Silicon Valley were associated with wiring the world, connecting individuals around the globe, and opening up access to vast stores of information, today they are just as likely to be known for doing the opposite.

Although Microsoft, Cisco, Yahoo, Skype, and Google have all come under scrutiny for colluding with China's internet

Table 23.1 Percent of tested websites blocked in 21 countries, by type of local and global content, 2006

Country	*Blocked websites, all*		*Blocked political websites*	
	Local	*Global*	*Local*	*Global*
Azerbaijan	2	98	0	100
Bahrain	62	38	86	14
China	80	20	92	8
Ethiopia	95	5	95	5
India	100	0	100	0
Iran	43	57	87	13
Libya	100	0	100	0
Myanmar	62	38	89	11
Oman	15	85	0	100
Pakistan	94	6	96	4
Saudi Arabia	22	78	68	32
Singapore	43	57	–	–
South Korea	81	19	100	0
Sudan	7	93	0	100
Syria	84	16	95	5
Thailand	81	19	100	0
Tunisia	32	68	73	27
United Arab Emirates	17	83	58	42
Uzbekistan	80	20	92	8
Vietnam	94	6	98	2
Yemen	10	90	22	78

Notes: Local websites are those designed for users in a specific country, and are usually in that country's national language. Global websites are primarily English language content, and include pornography.

censorship practices, perhaps the most significant, serious, and yet overlooked contribution to internet censorship by Western corporations comes from the manufacturers of the filtering software used to block content.

Internet security companies, like Fortinet, Secure Computing, and Websense, create off-the-shelf filtering products that block access to categorized lists of websites. While these products are primarily marketed to businesses, they have been readily employed by censoring states like Tunisia (Secure Computing), Iran (Secure Computing), Myanmar (Fortinet), and Yemen (Websense) to block access to politically-sensitive content.

Just like businesses that do not want their employees to view gambling or sport sites on company time, these governments simply tick off those categories of websites they do not want their citizens to access, such as "advocacy groups" or "militancy and extremist groups"—two categories in Websense's database. The former is defined by Websense as "sites that promote change or reform in public policy, public opinion, social practice, economic activities, and relationships," while the latter is defined as "sites that offer information about or promote or are sponsored by groups advocating anti-government beliefs or action" (Websense, 2007). As the Tunisia example listed in the introduction illustrates, however, the block lists used by these companies can contain categorization errors leading to untended blockages of websites.

Digital deceit

One troubling trend is been the lack of accountability and transparency over internet content-filtering practices by states that censor. While there is certainly a legitimate debate to be had about the balance between a state's right to cultural sovereignty and the free flow of information raised by internet censorship, unfortunately most states do not allow such a debate to take place prior to filtering, and have been shown to be deceitful about the content they block and the filtering practices they employ.

There are accountability and transparency issues around the disclosure of filtering practices. Among states that filter, few are willing to admit the full scope and scale and precise nature of their filtering systems. For example, Saudi Arabia provides a substantial level of detail about their filtering practices in published reports, including an acknowledgment on the block page it sends back to users' computers, as does the UAE. Other countries are not so open, and some engage in deceptive practices. For example, in China when users make a request for a website that is banned, the request is blocked at the router level and an error message is sent back to the user's machine effectively penalizing that machine's IP address from making further http requests for a varying period of time. From the user's end, the penalization appears as a "time-out" error with no explanation (Clayton *et al.*, 2006). Tunisia uses the same commercial product as Saudi Arabia—Smartfilter—but alters the block page functionality of that program to deliver a false error indication to users. When users attempt to access blocked content, they receive a page that appears to be a "File not found" error page but is in fact a block page designed to deceive users. In Uzbekistan, block pages sent back to users explain a site is blocked because it contains pornography even though the sites blocked are not pornographic but political in nature. Additionally, some Uzbeki ISPs redirect requests for banned content to unrelated sites or sites that are disguised to appear like the original site but which third parties operate and which contain false or misleading information. As countries that

censor are generally sensitive about it being known that they block access to political information, they tend to be opaque and/or deceptive about their filtering practices. Only very rarely do states fully disclose their filtering behavior.

As outlined in Table 23.2, most countries lack transparency and accountability when it comes to processes around internet filtering practices. Very few openly acknowledge filtering at all. Concealed filtering reflects either efforts to conceal the fact that filtering is occurring or the failure to clearly indicate filtering when it is occurring. Decentralized filtering is any blocking that occurs at the sub-national level, although this study does not include filtering at the institutional level, e.g., cybercafés, universities, or businesses. Transparency considers the presence of concealed filtering, provisions to appeal or report instances of inappropriate blocking, and open acknowledgement of filtering policies. Consistency measures the variation in filtering within a country across different ISPs.

Table 23.2 also shows the variation in filtering practices among countries in terms of relative centralization of content filtering methods and consistency among ISPs within the country. Many states defer the implementation of internet content filtering to individual ISPs, some of whom do not fully comply with authorities or choose their own methods or software products to perform the filtering. The result is that accessibility to internet content within certain countries, such as Azerbaijan, Burma, Iran, and Vietnam, for example, can vary widely depending on the ISP to which a user connects. When combined with a low

Table 23.2 Centralized, decentralized, concealed, transparency, and consistency of website filtering in 22 countries, 2006

Country	*Centralized filtering at national level*	*Decentralized filtering at sub-national level*	*Filtering concealed from user*	*Transparency of filtering policies*	*Consistency across country ISPs*
Azerbaijan		×		low	low
Bahrain		×	×	low	high
Burma		×		medium	low
China	×	×	×	low	medium
Ethiopia	×		×	low	high
India		×		high	medium
Iran		×		medium	medium
Jordan				low	high
Libya				low	high
Oman	×			high	high
Pakistan	×	×	×	high	medium
Saudi Arabia	×			high	high
Singapore		×		high	high
South Korea		×		high	high
Sudan	×			high	high
Syria		×		medium	high
Thailand		×		medium	medium
Tunisia	×			low	high
United Arab Emirates	×			medium	low
Uzbekistan	×		×	low	high
Vietnam		×	×	low	low
Yemen		×		medium	high

Source: Faris and Villeneuve, 2006.

degree of transparency and accountability, such a lack of consistency can be a vexing experience for users within the country who are unaware of what content is being denied to them and experience different forms of censorship depending on the ISP to which they connect.

Blocking by computer network attack and DNS tampering

Rather than blocking access to a site, entire websites can be forced offline and essentially silenced by attacks that overwhelm the servers that host the websites. For example, during elections in Kyrgyzstan, several opposition newspapers came under simultaneous distributed denial of service attacks. The opposition websites were moved to a hosting service at the Citizen Lab in Toronto for analysis. The attacks were carried out by a hacker or group of hackers known as "shadow team" based in the Ukraine, and although no conclusive proof could be obtained, the Kyrgyz authorities cannot be ruled out as being responsible. In addition to the attacks on the opposition websites, other attacks temporarily suspended access to all websites on two Kyrgyz ISPs (Elcat and AsiaInfo) (ONI, 2005d).

The same pattern of disruption during election periods was observed in Belarus in April 2006. Although no evidence of state-directed filtering or sponsorship of denial of service attacks could be found, there were several suspicious events:

- 37 opposition and media websites were inaccessible from the state-owned Beltelecom network on March 19 (election day), although they were accessible within Belarus from a different ISP network as well as from the external control location;
- the internet was inaccessible to subscribers using Minsk Telephone access numbers on March 25 (the day of a major demonstration, where riot police were used to disperse and arrest protestors);
- the website of the main opposition candidate, Aleksandr Milinkevich, was "dead" on March 19 and experienced access issues on March 21–22, (the post-election protest period); and
- an opposition website (Charter 97) was only partially accessible between March 19–25.

The internet is likely to be targeted by subtle methods of information disruption that are not so easily tracked and traced as are more traditional forms of filtering and surveillance. Moreover, the participants in these contests over information space are likely to include more than just state authorities, such as NGOs and activists, who benefit politically (with the outside world) by being able to claim they are under attack just as much as authorities may benefit by having their information kept offline (ONI, 2006).

The trends towards offensive computer network attacks as methods of filtering are even more significant in the context of the role the U.S. military is playing in setting doctrinal examples and establishing norms of acceptable practices in areas like information warfare. The recently declassified "Information Operations Roadmap" makes it clear that the U.S. and its regional allies intend on taking the war on terrorism to the internet, using a variety of means ranging from taking down "illegal content" through to using the internet as a means to "deter, deny and destroy terrorist groups" (U.S. Department of Defense, 2003).[6] Such militarization of cyberspace could legitimize the type of denial of service actions that occurred in Kyrgyzstan and Belarus, and open up dynamics of

competitive state and non-state offensive activities aimed at bringing down the sources of online information through "active," offensive means. Certainly the lessons have not been lost on the Chinese and Russian militaries, which are also supportive of a free-ranging scope for military action over the internet. Taken together with the shift in U.S. strategic policy towards preemption of threats "before they are fully formed," this stance has effectively opened the door for states to use computer network operations as a means to act unilaterally and extra-territorially to combat self-defined threats to national security emanating from the internet. As a consequence, computer network operations and information warfare are amongst the most secretive and fastest growing areas of investment for military, security, and signals intelligence organizations worldwide. Moreover, as the recent revelation concerning the U.S. National Security Agency's extralegal tapping of domestic communications (including the internet) suggest, even open and democratic societies are undertaking covert internet surveillance. The impact that these doctrinal shifts will have on the internet environment is likely to be substantial, and will make the challenges around accountability and transparency even more substantial.

Conclusion

Over the last several decades, the internet has enabled new, nimble, and distributed challenges to states worldwide, manifest in vigorous, mobilized opposition movements, protests, and in some cases, even revolutionary changes to political authority. Although these challenges have presented the most serious problems for non-democratic and authoritarian regimes, even among democratic states, the internet has presented serious challenges insofar as it empowers militant, terrorist, and criminal networks. Whereas once the promotion of new information and communication technologies were widely considered benign public policy, today states of all stripes have been pressed to find ways to limit and control the internet as a way to check their unintended and perceived negative consequences.

As the research shows, these efforts to control internet content are growing in scope, scale, and sophistication worldwide. Moreover, the methods used by states to filter content demonstrate a systematic lack of accountability and transparency. Although at first glance these policies and practices may be attributed simply to the strategic interests of states to control information flows across their territorial borders, the policies and practices of internet content filtering—in particular the use of computer network attacks and offensive information warfare—suggest a much deeper geopolitical struggle over the internet's architecture that is only beginning to unfold. Just as the domains of land, sea, air, and space have all been gradually colonized, militarized, and subject to inter-state competition so too is the once relatively unencumbered domain of cyberspace.

Of course these efforts by states to intervene in global internet communication flows are not going uncontested. The growth of state content-filtering practices has generated a burgeoning grass-roots transnational social movement around the protection and preservation of the internet as an open commons of information (see Deibert, 2003; Deibert and Rohozinski, 2007). The movement includes major NGOs, such as Amnesty International and Reporters without Borders, and efforts directed at multiple levels, from the construction of censorship circumvention technologies and other "hacktivist" tools to lobbying for the promotion of norms of openness and access to information at international levels.

These developments should make scholars of world politics and the internet rethink assumptions about not only the character of the internet but the social and political implications that flow from it. Although it is true that the internet helped unleash non-territorial forces and flows that have helped redefine the landscape of global politics, the internet's architecture is now being hotly contested and an object of competing discourses and practices of securitization. Almost certainly a new set of implications, many of them unintended, will flow as its architecture undergoes political transformation as a result of this competition.

Guide to further reading

In light of the fact that it is such a recent issue, there is relatively little scholarship about internet censorship and content filtering practices (outside of the work of the OpenNet Initiative outlined in this chapter). The latter is covered comprehensively in Deibert and Rohozinski (2007) with overviews of 41 countries and 8 regions, as well as several analytical chapters on the legal, social, and political implications of internet filtering. Those interested in exploring general issues of state control of internet communications might begin with Deibert (2003), Drezner (2004), as well as Goldsmith and Wu (2006), Goldsmith (1998), Lewis (2006), and Kalathil and Boas (2003). Villeneuve (2006) and Wu (2006) deal with some of the general issues concerning internet content filtering. There is a growing scholarship on internet content filtering in specific country and regional contexts, including Turkey (Altintos, 2002), China (Dowell, 2006; Hachigian, 2001; Lacharite, 2002; Li, 2003; Li, 2004), Singapore (Rodan, 1998), the Middle East (Goldstein, 1999), and Iran (Granick, 2005). Human Rights Watch (2006) did a major study on corporate complicity in internet censorship practices in China. Those interested in exploring some of the topics raised in this chapter concerning information warfare practices will find a much larger set of studies. Arquilla (1995, 1996), and Arquilla and Ronfeldt (2001) are essential background, with Adams (2001), Berkowitz (2003), Cohen (1996), Denning (1999), Der Derian (2000), Libicki (1998), Nye and Owens (1996), and Rattray (2001) all highly recommended as well.

Notes

1 OpenNet Initiative. http://opennet.net/
2 For additional detail on the analysis presented here, see Ronald J. Deibert, John G. Palfrey, Rafal Rohozinski and Jonathan Zittrain, (eds.) *Access Denied: the practice and policy of global internet filtering*, (MIT Press: 2008).
3 Those familiar with intelligence practices will recognize the combination of human and technical intelligence methods. The adoption of these methods, as well as other aspects by which the ONI operates, from intelligence-derived approaches has been deliberate. The ONI's researchers take considerable personal risks to carry out the tests, and great effort is taken to minimize those risks by securing group communications and employing compartmentalized information techniques. The latter means that researchers and experts consulted about or in a particular country may not know the identity of the testers in that country, and vise versa. Testing in some countries has been hampered by personal security considerations. In several instances, ONI researchers have been apprehended and interrogated by authorities for their activities.
4 Multiple categories of websites have been subject to internet filtering, including websites on a range of topics, from a range of site producers, or offering a range of services: free expression and media freedom; political transformation and opposition parties; political reform, legal reform, and governance; militants, extremists and separatists; human rights; foreign relations and military; minority rights and ethnic content; women's rights; environmental issues; economic development; sensitive or controversial history, arts and literature; hate speech; sex education and

family planning; public health; gay and lesbian content; pornography; provocative attire; dating; gambling; gaming; alcohol and drugs; minority faiths; religious conversion, commentary and criticism; anonymizers and circumvention; hacking; blogging domains and blogging services; web hosting sites and portals; voice over internet protocol (VOIP); free e-mail; search engines; translation; multimedia sharing; peer-to-peer file sharing; groups and social networking; commercial sites.

5 We group these categories themselves into four major categories: political, social, conflict/security, and internet tools.

6 The Pentagon document was written in October 2003, but recently obtained and released by a Freedom of Information request from the National Security Archives at George Washington University. It can be obtained from: http://www.gwu.edu/~nsarchiv/NSAEBB/NSAEBB177/info_ops_roadmap.pdf