

1

Measuring Global Internet Filtering

Robert Faris and Nart Villeneuve

The Scope and Depth of Global Internet Filtering

In this chapter, we set out to provide an overview of the data regarding Internet filtering that the OpenNet Initiative¹ has gathered over the past year. Empirical testing for Internet blocking was carried out in forty countries in 2006. Of these forty countries, we found evidence of technical filtering in twenty-six (see table 1.1). This does not imply that only these countries filter the Internet. The testing we carried out in 2006 constitutes the first step toward a comprehensive global assessment. Not only do we expect to find more countries that filter the Internet as we expand our testing, but we also expect that some of the countries that did not show signs of filtering in 2006 will institute filtering in subsequent years.²

Conceptually, the methodology we employ is simple. We start by compiling lists of Web sites that cover a wide range of topics targeted by Internet filtering. The topics are organized into a taxonomy of categories that have been subject to blocking, ranging from gambling, pornography, and crude humor to political satire and Web sites that document human rights abuses and corruption. (See table 1.2.) Researchers then test these lists to see which Web sites are available from different locations within each country.³

The states that filter the Internet must choose which topics to block (the scope of filtering) and how much of each topic to filter (the depth of filtering). The results of these decisions are summarized in figure 1.1, comparing the breadth and depth of filtering for the countries where evidence of filtering was found.

The number of different categories in which Internet filtering was found to occur is shown on the horizontal axis. We put this forward as a measure of the scope of Internet filtering in each country. (The categories are shown in table 1.2.)

The vertical axis depicts the comprehensiveness of filtering efforts as measured by the highest degree of content blocked in any of the topical categories. This captures a markedly different angle on filtering. If the breadth of filtering represents the ambition of censors to limit information related to a range of topics, the depth of filtering measures the success in actually blocking content. This might correspond to the level of sophistication of the filtering regime

and amount of resources devoted to the endeavor, or it may be a reflection of the resolve and political will to shut down large sections of the Internet.

The countries occupying the upper right of figure 1.1, including Iran, China, and Saudi Arabia, are those that not only intercede on a wide range of topics but also block a large amount of content relating to those topics. Myanmar and Yemen cover a similarly broad scope, though with less comprehensiveness in each category. South Korea is in a league of its own. It has opted to filter very little, targeting North Korean sites, many of which are hosted in Japan. Yet South Korea's thoroughness in blocking these sites manifests a strong desire to eliminate access to them. There is a cluster of states occupying the center of the plot that

Table 1.1

Filtering by state

Evidence of filtering	Suspected filtering	No evidence of filtering
Azerbaijan	Belarus	Afghanistan
Bahrain	Kazakhstan	Algeria
China		Egypt
Ethiopia		Iraq
India		Israel
Iran		Kyrgyzstan
Jordan		Malaysia
Libya		Moldova
Morocco		Nepal
Myanmar		Russia*
Oman		Ukraine
Pakistan		Venezuela
Saudi Arabia		West Bank/Gaza
Singapore		Zimbabwe
South Korea		
Sudan		
Syria		
Tajikistan		
Thailand		
Tunisia		
United Arab Emirates		
Uzbekistan		
Vietnam		
Yemen		

* Testing in Russia was limited to a selection of ISPs in Moscow; these preliminary results may not extend beyond this sample.

Table 1.2Categories subject to Internet filtering

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/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
P2P
Groups and social networking
Commercial sites

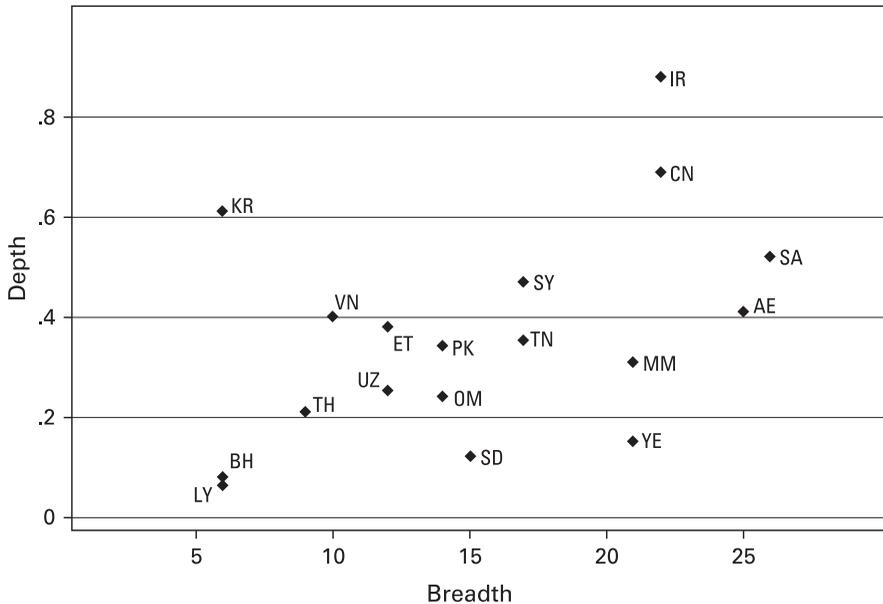


Figure 1.1

Comparing the breadth and depth of filtering. AE—United Arab Emirates; BH—Bahrain; CN—China; ET—Ethiopia; IR—Iran; JO—Jordan; KR—South Korea; LY—Libya; MM—Burma/Myanmar; OM—Oman; PK—Pakistan; SA—Saudi Arabia; SD—Sudan; SY—Syria; TH—Thailand; TH—Tunisia; UZ—Uzbekistan; VN—Vietnam; YE—Yemen. A number of countries that filter a small number of sites are omitted from this diagram, including Azerbaijan, Belarus, India, Jordan, Kazakhstan, Morocco, Singapore, and Tajikistan.

are widely known to practice filtering. These countries, which include Syria, Tunisia, Vietnam, Uzbekistan, Oman, and Pakistan, block an expansive range of topics with considerable depth. Ethiopia is a more recent entrant into this category, having extended its censorship of political opposition into cyberspace.

Azerbaijan, Jordan, Morocco, Singapore, and Tajikistan filter sparingly, in some cases as little as one Web site or a handful of sites. The evidence for Belarus and Kazakhstan remains inconclusive at the time of this writing, though blocking is suspected in these countries.

Of equal interest are the states included in testing in 2006 in which no evidence of filtering was uncovered (see table 1.1). We make no claims to have proven the absence of filtering in these countries. However, our background research supports the conclusion drawn from the technical testing that none of these states are currently filtering Internet content.⁴

Later in the book we turn our attention to the question of why some countries filter and others do not, even under similar political and cultural circumstances.

The Principal Motives and Targets of Filtering

On September 19, 2006, a military-led coup in Thailand overthrew the democratically elected government headed by Prime Minister Thaksin Shinawatra. Thailand is not unfamiliar with such upheavals. There have been seventeen coups in the past sixty years. This time, however, Internet users noticed a marked increase in the number of Web sites that were not accessible, including several sites critical of the military coup.⁵ A year earlier in Nepal, the king shut down the Internet along with international telephone lines and cellular communication networks when he seized power from the parliament and prime minister. In Bahrain, during the run-up to the fall 2006 election, the government chose to block access to a number of key opposition sites. These events are part of a growing global trend. Claiming control of the Internet has become an essential element in any government strategy to rein in dissent—the twenty-first century parallel to taking over television and radio stations.

In contrast to these exceptional events, the constant blocking of a swath of the Internet has become part of the everyday political and cultural reality of many states. A growing number of countries are blocking access to pornography, led by a handful of states in the Persian Gulf region. Other countries, including South Korea and Pakistan, block Web sites that are perceived as a threat to national security.

Notwithstanding the wide range of topics filtered around the world, there are essentially three motives or rationales for Internet filtering: politics and power, social norms and morals, and security concerns. Accordingly, most of the topics subject to filtering (see table 1.2) fall under one of three thematic headings: political, social, and security. A fourth theme—Internet tools—encompasses the networking tools and applications that allow the sharing of information relating to the first three themes. Included here are translation tools, anonymizers, blogging services, and other Web-based applications categorized in table 1.2.

Protecting intellectual property rights is another important driver of Internet content regulation, particularly in western Europe and North America. However, in the forty countries that were tested in 2006, this is not a major objective of filtering.⁶

Figure 1.2 compares the political and social filtering practices of these same twenty-seven countries. On one extreme is Saudi Arabia, which heavily censors social content. While there is also substantial political filtering carried out in Saudi Arabia, it is done with less scope and depth. On the other fringe are Syria and China, focusing much more of their extensive filtering on political topics. Myanmar and Vietnam are also notable for their primary focus on political issues, which in the case of Vietnam contradicts the stated reason for filtering the Internet.⁷ Iran stands out for its pervasive filtering of both political and social material.

Filtering directed at political opposition to the ruling government is a common type of blocking that spans many countries. Politically motivated filtering is characteristic of authoritarian and repressive regimes. The list of countries that engage in substantial political blocking includes Bahrain, China, Libya, Iran, Myanmar, Pakistan, Saudi Arabia, Syria, Tunisia,

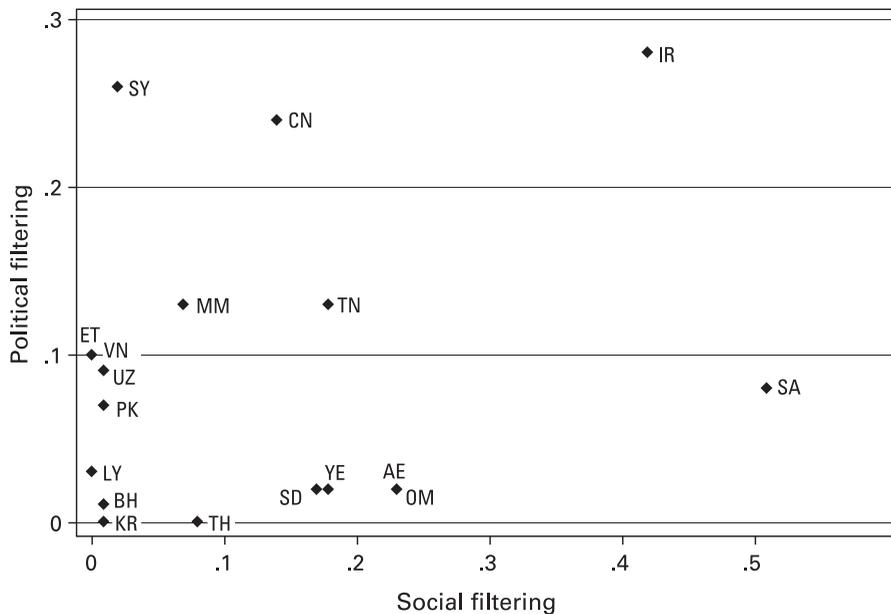


Figure 1.2

Political and social filtering. AE—United Arab Emirates; BH—Bahrain; CN—China; ET—Ethiopia; IR—Iran; JO—Jordan; KR—South Korea; LY—Libya; MM—Burma/Myanmar; OM—Oman; PK—Pakistan; SA—Saudi Arabia; SD—Sudan; SY—Syria; TH—Thailand; TH—Tunisia; UZ—Uzbekistan; VN—Vietnam; YE—Yemen. A number of countries that filter a small number of sites are omitted from this diagram, including Azerbaijan, Belarus, India, Jordan, Kazakhstan, Morocco, Singapore, and Tajikistan.

Uzbekistan, and Vietnam.⁸ Thailand and Ethiopia are the most recent additions to this group of countries that filter Web sites associated with political opposition groups. Yet in other countries with an authoritarian bent, such as Russia and Algeria, we have not uncovered filtering of the Internet.

The perceived threat to national security is a common rationale used for blocking content. Internet filtering that targets the Web sites of insurgents, extremists, terrorists, and other threats generally garners wide public support. This is best typified by South Korea where pro-North Korean sites are blocked, or by India where militant and extremist sites associated with groups that foment domestic conflict are censored. In Pakistan, Web sites devoted to the Balochi independence movement are blocked. Similarly, the Web sites of separatist or radical groups such as the Muslim Brotherhood are blocked in some countries in the Middle East.

Social filtering is focused on those topics that are held to be antithetical to accepted societal norms. Pornographic, gay and lesbian, and gambling-related content are prime examples

Box 1.1

Identifying and documenting Internet filtering

Measuring and describing Internet filtering defies simple metrics. Ideally, we would like to know how Internet censorship reduces the availability of information, how it hampers the development of online communities, and how it inhibits the ability of civic groups to monitor and report on the activities of the government, as these answers impact governance and ultimately economic growth. However, this is much easier to conceptualize at an abstract level than to measure empirically. Even if we were able to identify all the Web sites that have been put out of reach due to government action, the impact of blocking access to each Web site is far from obvious, particularly in this networked world where information has a habit of propagating itself and reappearing in multiple locations. Nevertheless, every obstacle thrown into the path of citizens seeking out information bears a cost or, depending on how one views the contribution of a particular Web site to society, a benefit. With this recognition of the inherent complexity of evaluating Internet censorship, we set out with modest goals—to identify and document filtering.

Two lists of Web sites are checked in each of the countries tested: a global list and a local list. The global list is a standardized list of Web sites that cover the categories listed in table 1.1. The global list of Web sites is comprised principally of internationally relevant Web sites with English content. The same global list is checked in each of the countries in which we have tested. A separate local list is created for each of the countries tested; it includes Web sites related to the specific issues and context of the study country.

These testing lists encompass a wide variety of content including political topics such as human rights, political commentary and news, religion, health and sex education, and Web sites sponsored by separatists and militant organizations. Pornography, gambling, drugs, and alcohol are also represented in the testing lists. The lists embody portions of the Web space that would be subject to Internet filtering in each of the countries being tested. They are designed to unearth filtering and blocking behavior where it exists. Background research is focused on finding sites that are likely to be blocked. In countries where Internet censorship has been reported, the lists include those sites that were alleged to have been blocked. These are not intended to be exhaustive lists of the relevant subject matter, nor do we presume to have identified all the Web sites that are subject to blocking.

The actual tests are run from within each country using software specifically designed for this purpose. Where appropriate, the tests are run from different locations to capture the differences in blocking behavior across Internet service providers (ISPs). The tests are run across multiple days and weeks to control for normal connectivity problems.

The completion of the initial accessibility testing is just the first step in the evaluation process. Additional diagnostic work is required to separate normal connectivity errors from intentional tampering. As described in further detail later, there are a number of technical alternatives for filtering the Internet, some of which are relatively easy to discover. Others are difficult to detect and require extensive diagnostic work to confirm.

of what is filtered for social and cultural reasons. Hate speech and political satire are also the target of Internet filtering in some countries. Web sites that deny the Holocaust or promote Nazism are blocked in France and Germany. Web sites that provide unflattering details related to the life of the king of Thailand are censored in his country.

An emergent impetus for filtering is the protection of existing economic interests. Perhaps the best example is the blocking of low-cost international telephone services that use Voice-over Internet Protocol (VoIP) and thereby reduce the customer base of large telecommunications companies, many of which enjoy entrenched monopoly positions. Skype, a popular and low-cost Internet-based telephone service, has been blocked in Myanmar and United Arab Emirates, which heavily block VoIP sites. The Web sites of many VoIP companies are also blocked in Syria and Vietnam.

Many countries seek to block the intermediaries: the tools and applications of the Internet that assist users in accessing sensitive material on the Internet. These tools include translation sites, e-mail providers, Weblog hosting sites, and Web sites that allow users to circumvent standard blocking strategies. Blogging services such as Blogspot are often targeted; eight countries blocked blogs hosted there, while Syria, Ethiopia, and Pakistan blocked the entire domain, denying access to all the blogs hosted on Blogspot. Fourteen countries blocked access to anonymity and censorship circumvention sites. Both SmartFilter, used in Sudan, Tunisia, Saudi Arabia, and UAE, and Websense, used in Yemen, have filtering categories—called “Anonymizers” and “Proxy Avoidance,” respectively—used to block such sites.

A handful of countries, including China, Vietnam, and states in the MENA region (the Middle East and North Africa), block Web sites related to religion and minority groups. In China, Web sites that represent the Falun Gong and the Tibetan exile groups are widely blocked. In Vietnam, religious and ethnic sites associated with Buddhism, the Cao Dai faith, and indigenous hill tribes are subject to blocking. Web sites that are aimed at religious conversion from Islam to Christianity are often blocked in the MENA region. Decisively identifying the motives of filtering activity is often impossible, particularly as the impact of filtering can simultaneously touch a host of social and political processes. That being said, it probably would be a mistake to attribute the filtering of religious and ethnic content solely to biases against minority groups, as these movements also represent a political threat to the ruling regimes.

A Survey of Global Filtering Strategies, Transparency, and Consistency

There are many techniques used to block access to Internet content. Each of these techniques can be used at different levels of Internet access within a country. Internet filtering is most commonly implemented at two levels: at the ISPs within the country and on the Internet backbone at the international gateway. These methods may overlap; an ISP may filter content using one particular technique while another technique is used at the international gateway.

Pakistan is an example of a country that blocks at both the international gateway and at the ISP level.

There are a few principal techniques used for Internet filtering including IP blocking, DNS tampering, and proxy-based blocking methods. (For blocking behavior by country, see table 1.3.) These techniques are presented in further detail by Anderson and Murdoch in chapter 3.

IP blocking is effective in blocking the intended target and no new equipment needs to be purchased. It can be implemented in an instant; all the required technology and expertise is

Table 1.3
Blocking techniques

	IP blocking	DNS tampering	Blockpage	Keyword
Azerbaijan	X		X	
Bahrain		X	X	
China	X			X
Ethiopia	X			
India	X	X		
Iran			X	X
Jordan	X			
Libya	X			
Myanmar			X	
Oman			X	
Pakistan	X	X		
Saudi Arabia			X	
Singapore			X	
South Korea	X	X	X	
Sudan			X	
Syria			X	
Thailand			X	
Tunisia			X	
United Arab Emirates			X	
Uzbekistan*			X	
Vietnam		X	X	
Yemen			X	X

Blocking behavior included in this table may include international gateway level filtering, and filtering techniques used by different ISPs.

* In Uzbekistan, the blockpage does not clearly indicate that filtering is occurring but rather redirects users to a third-party Web site.

readily available. Depending on the network infrastructure within the country it may also be possible to block at or near the international gateways so that the blocking is uniform across ISPs.

Countries new to filtering will generally start with IP blocking before moving on to more expensive filtering solutions. ISPs most often respond quickly and effectively to blocking orders from the government or national security and intelligence services. Therefore they block what is requested in the cheapest way using technology already integrated into their normal network environment. Blocking by IP can result in significant overblocking as all other (unrelated) Web sites hosted on that server will also be blocked.

China uses IP blocking to obstruct access to at least three hundred IP addresses. This blocking is done at the international gateway level affecting all users of the network regardless of ISP. The IPs blocked among the two backbone providers, China Netcom and ChinaTelecom, are remarkably similar.⁹

The ISP ETC-MC in Ethiopia uses IP blocking to block, among other sites, Google's Blogspot blogging service. This results in all Blogspot blogs being blocked in Ethiopia. Pakistan implements IP blocking at the international gateway level. In addition to blocking the IP for Blogspot, they also block Yahoo's hosting service, which results in major overblocking. For example, in targeting www.balochvoice.com they are actually blocking more than 52,000 other Web sites hosted on that same server.

DNS tampering is achieved by purposefully disrupting DNS servers, which resolve domain names into IP addresses. Generally, each ISP maintains its own DNS server for use by its customers. To block access to particular Web sites, the DNS servers are configured to return the wrong IP address. While this allows the blocking of specific domain names, it also can be easily circumvented by simple means such as accessing an IP address directly or by configuring your computer to use a different DNS server.

In Vietnam, the ISP FPT configures DNS to not resolve certain domain names, as if the site does not exist. The ISP Cybernet in Pakistan also uses this technique. The ISP Batelco in Bahrain uses this technique for some specific opposition sites. Batelco did not, however, completely remove the entry (the MX record for e-mail still remains). In India, the ISP BHARTI resolves blocked sites to the invalid IP address 0.0.0.0 while the ISP VSNL resolves blocked sites to the invalid IP address 1.2.3.4. The South Korean ISP, Hananet, uses this technique but makes the blocked Web site resolve to 127.0.0.1. This is the IP address for the "localhost." Another South Korean ISP, KORNET, makes blocked sites resolve to an ominous police Web site. This represents an unusual case in which DNS tampering resolves to a block-page.¹⁰

Our tests revealed that there is often a combination of IP blocking and DNS tampering. It may be a signal that countries are responding to the outcry concerning the overblocking associated with IP blocking and moving to the targeting of specific domain names with DNS tam-

pering. In India, for example, the Internet Service Providers Association of India reportedly has sent instructions to ISPs showing how to block by DNS instead of by IP.¹¹

In proxy-based filtering strategies, Internet traffic passing through the filtering system is reassembled and the specific HTTP address being accessed is checked against a list of blocked Web sites. These can be individual domains, subdomains, specific long URL paths, or keywords in the domain or URL path. When users attempt to access blocked content they are subsequently blocked. An option in this method of filtering is to return a *blockpage* that informs the user that the content requested has been blocked.

Saudi Arabia uses SmartFilter as a filtering proxy and displays a blockpage to users when they try to access a site on the country's block list. The blockpage also contains information on how to request that a block be lifted. Saudi Arabia blocks access to specific long URLs. For example, www.humum.net/ is accessible, while www.humum.net/country/saudi.shtml is blocked. United Arab Emirates, Oman, Sudan, and Tunisia also use SmartFilter. Tunisia uses SmartFilter as a proxy to filter the Internet. But instead of showing users a blockpage indicating that the site has been blocked, they have created a blockpage that looks like the Internet Explorer browser's default error page (in French), presumably to disguise the fact that they are blocking Web sites.

A proxy-based filtering system can also be programmed such that Internet traffic passing through the filtering system is reassembled and the specific HTTP address requested is checked against a list of blocked keywords. No country that ONI tested blocked access to a Web site as a result of a keyword appearing in the body content of the page, however, there are a number of countries that block by keyword in the domain or URL path, including China, Iran, and Yemen.

China filters by keywords that appear in the host header (domain name) or URL path. For example, while the site <http://archives.cnn.com/> is accessible, the URL <http://archives.cnn.com/2001/ASIANOW/east/01/11/falun.gong.factbox/> is not. When this URL is requested, reset (RST) packets are sent that disrupt the connection, presumably because of the keyword *falun.gong*. Iran uses a filtering proxy that displays a blockpage when a blocked Web site is requested. On some ISPs in Iran, such as Shatel and Datak, keywords in URL paths are blocked. This most often affects search queries in search engines. For example, here is a query run on Google for *naked* in Arabic (www.google.com/search?hl=fa&q=%D9%84%D8%AE%D8%AA&btnG=%D8%A8%D9%8A%D8%A7%D8%A8) that was blocked. Ynet in Yemen blocks any URL containing the word *sex*. The domain www.arabtimes.com is blocked in Oman and the UAE but the URL for the Google cached version (<http://72.14.235.104/search?q=cache:8utpDVLa1yYJ:www.arabtimes.com/+arabtimes&hl=en&ct=clnk&cd=1>) is also blocked because *www.arabtimes.com* appears in the URL path.

Filtering systems can also be configured to redirect users to another Web site. In most cases, redirection is identical to blockpage filtering, the only difference being the route used

to produce the blockpage. ISPs in Iran, Singapore, Thailand, and Yemen all use redirection to a blockpage. Uzbekistan uses redirection but instead of redirecting to a blockpage the filters send users to Microsoft's search engine at www.live.com, suggesting that the government wishes to conceal that fact that blocking has taken place.

There are thus various degrees of transparency in Internet filtering. Where blockpages are used, it is clearly apparent to users when a requested Web site has been intentionally blocked. Other countries give no indication that a Web site is blocked. In some cases, this is a function of the blocking technique being used. Some countries, such as Tunisia and Uzbekistan, appear to deliberately disguise the fact that they are filtering Internet content, going a step farther to conceal filtering activity beyond the failure to inform users that they are being filtered.

Another subset of countries, including Bahrain and United Arab Emirates, employ a hybrid strategy, indicating clearly to users that certain sites are blocked while obscuring the blocking of other sites behind the uncertainty of connection errors that could have numerous other explanations. In Bahrain, users normally receive a blockpage. However, for the specific site www.vob.org, Bahrain uses DNS tampering that results in an error. In United Arab Emirates all blocked sites with the exception of www.skype.com returned a blockpage. There is an apparent two-tiered system in place. They are willing to go on the record as blocking some sites, and not for others.

Providing a blockpage informing a user that their choice of Web site is not available by action of the government is still short of providing a rationale for the blocking of that particular site, or providing a means for appealing this decision. Very few countries go this far. A small group of countries, including Saudi Arabia, Oman, and United Arab Emirates, and some ISPs in Iran, allow Internet users to write to authorities to register a complaint that a given Web site has been blocked erroneously.

Centralized filtering regimes require all Internet traffic to pass through the same filters. This results in a consistent view of the Internet for users within the country; all users experience the same degree of filtering. This is most commonly implemented at the international gateway. When filtering is delegated to the ISP level, and hence decentralized, there may be significant differences among ISPs regarding the filtering techniques used and the content that is filtered. In this case, access to Web sites may vary substantially depending on the blocking choices of individual ISPs. (Table 1.4 presents the use of centralized and/or decentralized filtering strategies across the countries in the study, and the resulting consistency in filtering within each country.) In Iran there is considerable variation in the blocking among ISPs. For example, one ISP blocks considerably less political content than the other six ISPs tested. Only one ISP out of the five tested in Azerbaijan, AzNet, blocks access to a considerable amount of social content, most of which is pornographic, while the others block access to only a single IP address. In Myanmar, there is substantial variation in the filtering between the two ISPs tested. One filters much more pornography, while the other blocks a significantly greater portion of politically oriented Web sites. In the United Arab Emirates, an ISP that serves primarily the free-trade

Table 1.4
Comparing filtering regimes

	Locus	Consistency	Concealed filtering	Transparency and accountability
Azerbaijan	D	Low		Medium
Bahrain	C	High	Yes	Low
China	C and D	Medium	Yes	Low
Ethiopia	C	High	Yes	Low
India	D	Medium		High
Iran	D	Medium		Medium
Jordan	D	High		Low
Libya	C	High	Yes	Low
Morocco	C	High	Yes	Low
Myanmar	D	Low		Medium
Oman	C	High		High
Pakistan	C and D	Medium	Yes	High
Saudi Arabia	C	High		High
Singapore	D	High		High
South Korea	D	High		High
Sudan	C	High		High
Syria	D	High		Medium
Tajikistan	D	Low		Medium
Thailand	D	Medium		Medium
Tunisia	C	High	Yes	Low
United Arab Emirates	D	Low		Medium
Uzbekistan	C and D	High	Yes	Low
Vietnam	D	Low	Yes	Low
Yemen	D	High		Medium

The **Locus** of filtering indicates where Internet traffic is blocked. **C** indicates that traffic is blocked from a central location, normally the Internet backbone, and affects the entire state equally. **D** indicates that blocking is decentralized, typically implemented by ISPs. (Note that this study does not include filtering at the institutional level, for example, cybercafés, universities, or businesses.)

Consistency measures the variation in filtering within a country across different ISPs where applicable.

Concealed filtering reflects either efforts to conceal the fact that filtering is occurring or the failure to clearly indicate filtering when it occurs.

Transparency and accountability corresponds to the overall level of openness in regard to the practice of filtering. It also considers the presence of concealed filtering, the type of notice given to users regarding blocking, provisions to appeal or report instances of inappropriate blocking, and public acknowledgement of filtering policies.

zone has not historically filtered the Internet, while the predominant ISP for the rest of the country has consistently filtered the Internet.

Modifications can be made to the blocking efforts of a country by the authorities at any time. Sites can be added or removed at their discretion. For example, during our tests in Iran the Web site of the *New York Times* was blocked, but for only one day. Some countries have also been suspected of introducing temporary filtering around key time periods such as elections.

Hosting modifications can also be made to a blocked site resulting in it becoming accessible or inaccessible. For example, while Blogspot blogs were blocked in Pakistan due to IP blocking, the interface to update one's blog was still accessible. However, Blogspot has since upgraded its service and the new interface is hosted on the blocked IP, making the interface to update one's blog inaccessible in Pakistan. The reverse is also possible. For example, if the IP address of a Web site is blocked, the Web site may change its hosting arrangement in order to receive a new IP address, leaving it unblocked until the new IP address is discovered and blocked.

Summary Measures of Internet Filtering

To summarize the results of our work, we have assigned a score to each of the countries we studied. This score is designed to reflect the degree of filtering in each of the four major thematic areas: 1) the filtering of political content, 2) social content, 3) conflict- and security-related content, and 4) Internet tools and applications. Each country is given a score on a four-point scale that captures both the breadth and depth of filtering for content of each thematic type (see table 1.5).

- Pervasive filtering is defined as blocking that spans a number of categories while blocking access to a large portion of related content.
- Substantial filtering is assigned where either a number of categories are subject to a medium level of filtering in at least a few categories or a low level of filtering is carried out across many categories.
- Selective filtering is either narrowly defined filtering that blocks a small number of specific sites across a few categories, or filtering that targets a single category or issue.
- Suspected filtering is assigned where there is information that suggests that filtering is occurring, but we are unable to conclusively confirm that inaccessible Web sites are the result of deliberate tampering.

The scores in table 1.5 are subjective evaluations based upon the quantitative information gathered during a year of testing and research. In 2006, we tested thousands of Web sites across more than 120 ISPs in 40 countries, creating a database with close to 200,000 observations. Each observation is in turn based on the conclusion of an average of ten accessibility tests. Despite the breadth of this data, a purely quantitative reporting might be

Table 1.5
Summary of filtering

	Political	Social	Conflict and security	Internet tools
Azerbaijan	●	—	—	—
Bahrain	●●	●	—	●
Belarus	○	○	—	—
China	●●●	●●	●●●	●●
Ethiopia	●●	●	●	●
India	—	—	●	●
Iran	●●●	●●●	●●	●●●
Jordan	●	—	—	—
Kazakhstan	○	—	—	—
Libya	●●	—	—	—
Morocco	—	—	●	●
Myanmar	●●●	●●	●●	●●
Oman	—	●●●	—	●●
Pakistan	●	●●	●●●	●
Saudi Arabia	●●	●●●	●	●●
Singapore	—	●	—	—
South Korea	—	●	●●●	—
Sudan	—	●●●	—	●●
Syria	●●●	●	●	●●
Tajikistan	●	—	—	—
Thailand	●	●●	—	●
Tunisia	●●●	●●●	●	●●
United Arab Emirates	●	●●●	●	●●
Uzbekistan	●●	●	—	●
Vietnam	●●●	●	—	●●
Yemen	●	●●●	●	●●

●●● Pervasive filtering; ●● Substantial filtering; ● Selective filtering; ○ Suspected filtering; — No evidence of filtering.

misleading unless we were able to effectively measure the relative importance of each Web site. For example, the blocking of BBC or Wikipedia represents far more than the blocking of a less prominent Web site. Similarly, blocking a social networking site or a blogging server could have a profound impact on the formation of online communities and on the publication of user-generated content. While Internet users will eventually provide alternatives to recreate these communities on other sites hosted on servers that are not blocked, the transition of a wide community is unlikely given the time, effort, and coordination required to reconstitute a community in another location. At the other extreme, the blocking of one pornographic site will have a minor impact on Internet life if access to thousands of similar sites remains unimpeded. For these reasons, we have decided to summarize the results of testing categorically, considering both the scope and depth of the quantitative testing results, in conjunction with expert opinion regarding the significance of the blocking of individual Web sites.

It is tempting to aggregate the results by summing up the scores in each category. Yet doing so would imply that the blocking of political opposition is equivalent to filtering that supports conservative social values or the fear of national security risks. These competing sets of values suggest that a number of different weighting schemes might be appropriate. In any case, the results are generally quite clear, as the most pervasive filtering regimes tend to filter across all categories.

Country-specific and Global Filtering

A comparison between the blocking of country-specific sites and the blocking of internationally relevant Web sites provides another view of global filtering. Not surprisingly, we found that

Box 1.2

Where we tested

The decision where to test was a simple pragmatic one—where were we able to safely test and where did we have the most to learn? Two countries did not make the list this year because of security concerns: North Korea and Cuba. Learning more about the filtering practices in these countries is certainly of great interest to us. However, we were not confident that we could adequately mitigate the risks to those who would collaborate with us in these countries.

A number of other countries in Europe and North America that are known to engage in filtering to varying degrees were not tested this year. This decision again was a fairly easy practical choice. The filtering practices in these countries are better understood than in other parts of the world and we therefore had less to contribute here. Many of the countries in Europe focus their Internet filtering activity on child pornography. This is not a topic that we will test for ethical and legal reasons.

the incidence of blocking Web sites in our testing lists was approximately twice as high for Web sites available in a local language compared to sites available only in English or other international languages. Figure 1.3 shows that many countries focus their efforts on filtering locally relevant Web content. Ethiopia, Pakistan, Syria, Uzbekistan, and Vietnam are examples of countries that extensively block local content while blocking relatively few international Web sites. China and Myanmar also concentrate more of their filtering efforts on country-specific Internet content, though they block somewhat more global content. Middle Eastern filtering regimes tend to augment local filtering with considerably more global content. This balance mirrors the use of commercial software, generally developed in the West, to identify and block Internet content.

Table 1.6 shows an alternative view of filtering behavior, looking at the blocking of different types of content providers rather than content. The apparent prime targets of filtering are blogs, political parties, local NGOs, and individuals. In the case of blogs, a number

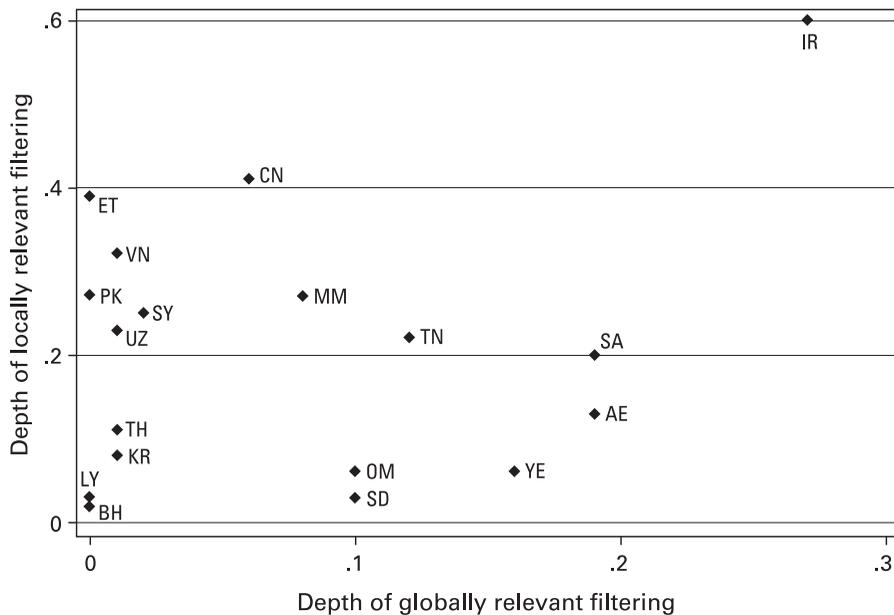


Figure 1.3

Filtering targeted at local sites and global sites. AE—United Arab Emirates; BH—Bahrain; CN—China; ET—Ethiopia; IR—Iran; JO—Jordan; KR—South Korea; LY—Libya; MM—Burma/Myanmar; OM—Oman; PK—Pakistan; SA—Saudi Arabia; SD—Sudan; SY—Syria; TH—Thailand; TH—Tunisia; UZ—Uzbekistan; VN—Vietnam; YE—Yemen. A number of countries that filter a small number of sites are omitted from this diagram, including Azerbaijan, Belarus, India, Jordan, Kazakhstan, Morocco, Singapore, and Tajikistan.

of countries, including Pakistan and Ethiopia, have blocked entire blogging domains, which inflates these figures. Logically, these assessments represent more accurately the result of filtering rather than the intention. Establishing the intention of blocking is never as clear. The blocking of this wide array of blogs could be the result of a lack of technical sophistication or a desire to simultaneously silence the entire collection of blogs hosted on the site.

The other prominent target of filtering is political parties, followed by NGOs focused on a particular region or country, and Web sites run by individuals. The implications of targeting civic groups and individual blogs are addressed by Deibert and Rohozinski in chapter 6 of this volume.

First Steps Toward Understanding Internet Filtering

In this chapter, we summarize what we have learned over the past year regarding the incidence of global Internet filtering. Taking an inventory of filtering practices and strategies is a necessary and logical first step, though still far from a thorough understanding of the issue. The study of Internet filtering can be approached by asking why some states filter the Internet or by asking why others do not. The latter question is particularly apt in countries that maintain a repressive general media environment while leaving the Internet relatively open. This is not

Table 1.6

Blocking by content provider

Content provider type	Portion of content filtered
Academic	0.02
Blogs	0.20
Chat and discussion boards	0.05
Government	0.03
Government media	0.02
International governmental organizations	0.00
Independent media	0.06
Individual	0.09
International NGOs	0.02
Labor groups	0.05
Locally focused NGOs	0.09
Militant groups	0.01
Political parties	0.19
Private businesses	0.06
Religious groups	0.02
Regional NGOs	0.04

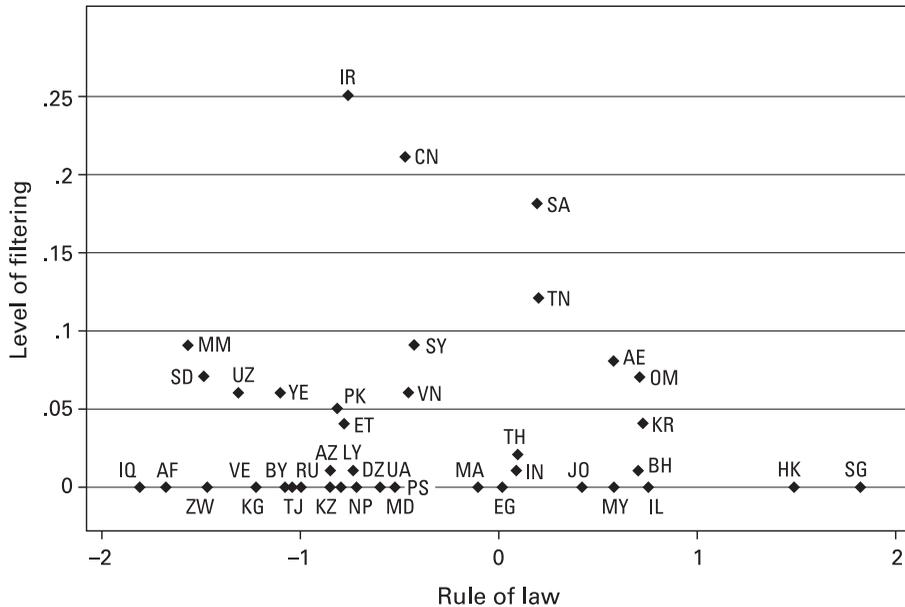


Figure 1.4

Filtering and the rule of law. AE—United Arab Emirates; AF—Afghanistan; AZ—Azerbaijan; BH—Bahrain; BY—Belarus; CN—China; DZ—Algeria; EG—Egypt; ET—Ethiopia; HK—Hong Kong; IL—Israel; IN—India; IR—Iran; IQ—Iraq; JO—Jordan; KG—Kyrgyzstan; KR—South Korea; KZ—Kazakhstan; LY—Libya; MA—Morocco; MD—Moldova; MM—Burma/Myanmar; MY—Malaysia; NP—Nepal; OM—Oman; PK—Pakistan; PS—Gaza/West Bank; RU—Russia; SA—Saudi Arabia; SD—Sudan; SG—Singapore; SY—Syria; TH—Thailand; TH—Tunisia; TN—Tunisia; TJ—Tajikistan; UA—Ukraine; UZ—Uzbekistan; VE—Venezuela; VN—Vietnam; YE—Yemen; ZW—Zimbabwe.

an uncommon circumstance. Pointing simply toward the absence of a solid rule of law does not seem promising. As seen in figure 1.4, there is no simple relationship between the rule of law and filtering, at least not as rule of law is defined and measured by the World Bank.¹² A country can maintain a better-than-average rule of law record and still filter the Internet. Similarly, many countries suffer from a substandard legal situation while maintaining an open Internet.

Comparing filtering practices with measures of voice and accountability is more telling. The countries that actively engage in the substantial filtering of political content also score poorly on measures of voice and accountability. This is true for both political and social Internet blocking, as shown in figures 1.5 and 1.6. Yet many of the anomalies persist. We are still far from explaining why some countries resort to filtering while others refrain from taking this step. This does stress the diversity of strategies and approaches that are being taken to regulate

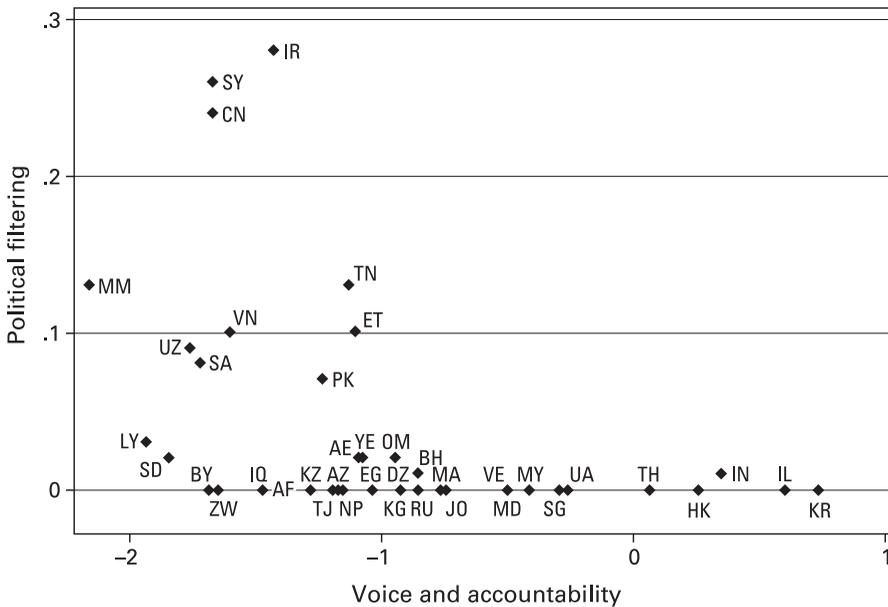


Figure 1.5

Political filtering and voice and accountability. AE—United Arab Emirates; AF—Afghanistan; AZ—Azerbaijan; BH—Bahrain; BY—Belarus; CN—China; DZ—Algeria; EG—Egypt; ET—Ethiopia; HK—Hong Kong; IL—Israel; IN—India; IR—Iran; IQ—Iraq; JO—Jordan; KG—Kyrgyzstan; KR—South Korea; KZ—Kazakhstan; LY—Libya; MA—Morocco; MD—Moldova; MM—Burma/Myanmar; MY—Malaysia; NP—Nepal; OM—Oman; PK—Pakistan; PS—Gaza/West Bank; RU—Russia; SA—Saudi Arabia; SD—Sudan; SG—Singapore; SY—Syria; TH—Thailand; TH—Tunisia; TN—Tunisia; TJ—Tajikistan; UA—Ukraine; UZ—Uzbekistan; VE—Venezuela; VN—Vietnam; YE—Yemen; ZW—Zimbabwe.

the Internet. We are also observing a recent and tremendously dynamic process. The view we have now may change dramatically in the coming years.

The link between repressive regimes and political filtering follows a clear logic. However, the link between regimes that suppress free expression and social filtering activity is less obvious. Part of the answer may reside in that regimes that tend to filter political content also filter social content.

Figure 1.7 demonstrates that few states restrict their activities to one or two types of content. Once filtering is implemented, it is applied to a broad range of content. These different types of filtering activities are often correlated with each other, and can be used as a pretense for expanding government control of cyberspace.

Vietnam, for example, uses pornography as its publicly stated reason for filtering, yet blocks little pornography. It does, however, filter political Internet content that opposes one-party rule

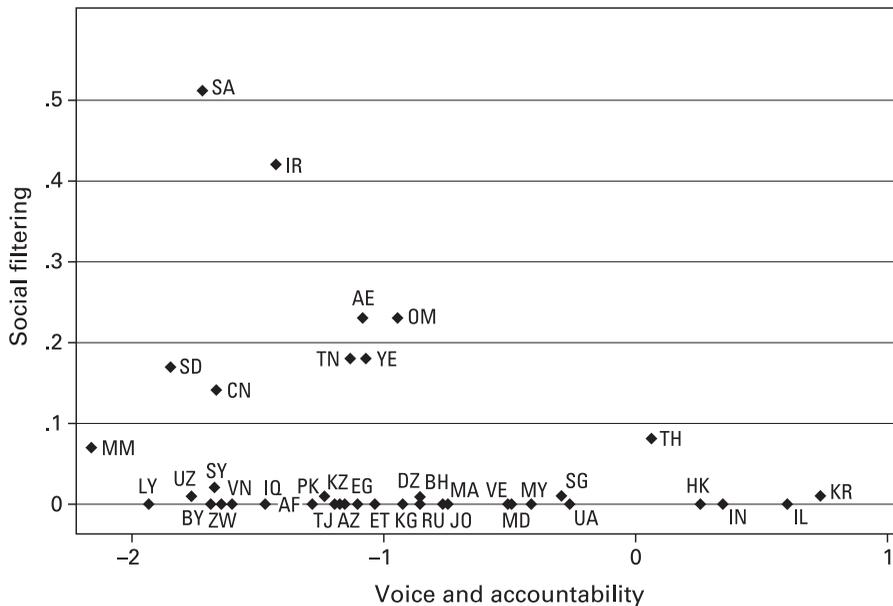


Figure 1.6

Social filtering and voice and accountability. AE—United Arab Emirates; AF—Afghanistan; AZ—Azerbaijan; BH—Bahrain; BY—Belarus; CN—China; DZ—Algeria; EG—Egypt; ET—Ethiopia; HK—Hong Kong; IL—Israel; IN—India; IR—Iran; IQ—Iraq; JO—Jordan; KG—Kyrgyzstan; KR—South Korea; KZ—Kazakhstan; LY—Libya; MA—Morocco; MD—Moldova; MM—Burma/Myanmar; MY—Malaysia; NP—Nepal; OM—Oman; PK—Pakistan; PS—Gaza/West Bank; RU—Russia; SA—Saudi Arabia; SD—Sudan; SG—Singapore; SY—Syria; TH—Thailand; TH—Tunisia; TN—Tunisia; TJ—Tajikistan; UA—Ukraine; UZ—Uzbekistan; VE—Venezuela; VN—Vietnam; YE—Yemen; ZW—Zimbabwe.

in Vietnam. In Saudi Arabia and Bahrain, filtering does not end with socially sensitive material such as pornography and gambling but expands into the political realm.

Once the technical and administrative mechanisms for blocking Internet content have been put into place, it is a trivial matter to expand the scope of Internet censorship. As discussed in subsequent chapters, the implementation of filtering is often carried by private sector actors—normally the ISPs—using software developed in the United States. Filtering decisions are thus often made by selecting categories for blocking within software applications, which may also contain categorization errors resulting in unintended blocking. The temptation and potential for mission creep is obvious. This slope is made ever more slippery by the fact that transparency and accountability are the exception in Internet filtering decisions, not the norm.

In the following chapter, Zittrain and Palfrey probe in further detail the political motives and implications of this growing global phenomenon, with subsequent chapters elaborating on technical, legal, and ethical considerations.

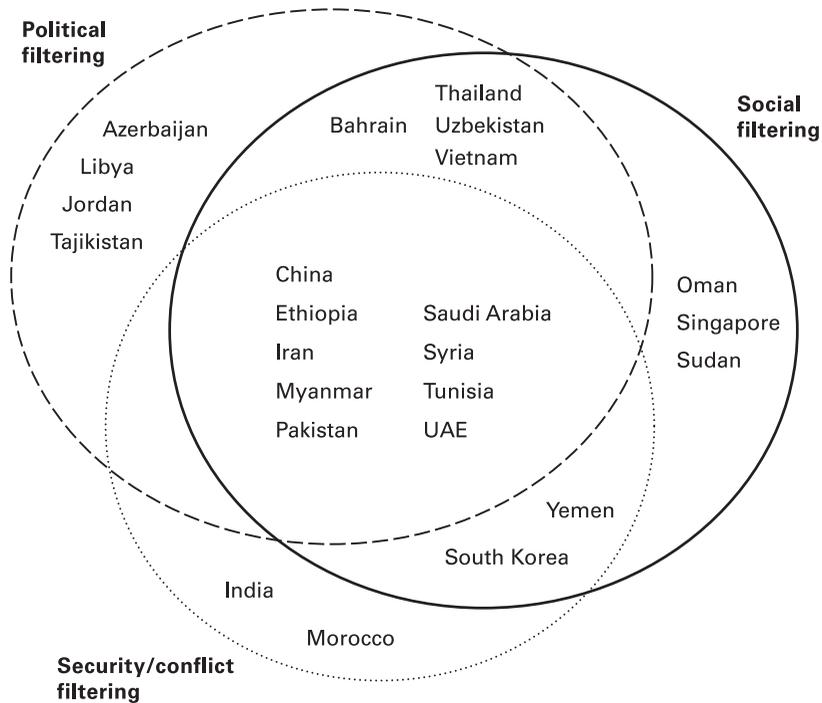


Figure 1.7
Content filtering choices.

Notes

1. The OpenNet Initiative is a collaboration of four institutions: the Citizen Lab at the University of Toronto, the Oxford Internet Institute at Oxford University, the Berkman Center for Internet & Society at Harvard Law School, and the University of Cambridge. More information is available at <http://www.opennetinitiative.net>.
2. A number of countries are currently debating strategies and legislation to filter the Internet, including Norway, Russia, and many countries in Latin America.
3. Each set of tests is performed on different Internet service providers within the country.
4. The Internet filtering tests carried out in Russia in 2006 were limited to ISPs accessible in Moscow. These results therefore do not necessarily reflect the situation in other areas of the country.
5. The blocking of two sites garnered most of the attention: one devoted to opposition to the September 19 coup (<http://www.19sep.com/>) and another hosted by Thai academics (<http://www.midnightuniv.org/>).
6. The strategies for addressing alleged intellectual property rights violations can vary significantly from standard Internet filtering. Rather than blocking Web sites that continue to be available from other locations, efforts generally focus on taking down the content from the Web sites that have posted the material and on removing the sites from the results of search engines. Moreover, takedown efforts are often instigated by private parties with the threat of subsequent legal action rather than being initiated by government action. See www.chillingeffects.org for more information.

-
7. The ONI Vietnam report is available at http://www.opennetinitiative.net/studies/vietnam/ONI_Vietnam_Country_Study.pdf.
 8. We were not able to test in Cuba or North Korea. Both countries are reported to engage in pervasive filtering in addition to curtailing access to the Internet. See "Going Online in Cuba: Internet under Surveillance," http://www.rsf.org/IMG/pdf/rapport_gb_md_1.pdf, and Tom Zeller, "The Internet Black Hole That Is North Korea," *New York Times*, 23 October 2006.
 9. There are two principal ISPs in China—one that covers the north and one the south. The smaller ISPs in China that serve Internet users connect to the Internet backbone through one of these large ISPs.
 10. It also demonstrates that the use of DNS tampering does not necessitate a lack of transparency in filtering. If it were deemed important, users could be informed that the Web site they were seeking was being intentionally blocked.
 11. See Shivam Vij, "Blog Blockade Will Be Lifted in 48 Hours," Rediff India Abroad, <http://www.rediff.com/news/2006/jul/19blogs.htm>.
 12. Information on the compilation and estimation of the "rule of law" and "voice and accountability" measures are available at the World Bank Governance and Anti-Corruption Web site, www.worldbank.org/wbi/governance. Their definitions of these indicators are: "Voice and Accountability includes in it a number of indicators measuring various aspects of the political process, civil liberties, political and human rights, measuring the extent to which citizens of a country are able to participate in the selection of governments." "Rule of Law includes several indicators which measure the extent to which agents have confidence in and abide by the rules of society. These include perceptions of the incidence of crime, the effectiveness and predictability of the judiciary, and the enforceability of contracts."

