The Evolving Landscape of Internet Control

Hal Roberts, Ethan Zuckerman, Robert Faris, Jillian York, and John Palfrey

A Summary of Our Recent Research and Recommendations

Berkman Center for Internet & Society

August 2011

Each co-author is affiliated with The Berkman Center for Internet & Society at Harvard University. Hal Roberts is a fellow; Ethan Zuckerman is a senior researcher; Robert Faris is the director of research; Jillian York was a Berkman Center staff member at the time of this report; and John Palfrey is a faculty co-director of the Berkman Center. We gratefully acknowledge support for this research from the US Department of State via a subgrant through Internews Network (USA).

OVERVIEW

Over the past two years, we have undertaken several studies at the Berkman Center designed to better understand the control of the Internet in less open societies. During the years we've been engaged in this research, we have seen many incidents that have highlighted the continuing role that the Internet serves as a battleground for political control, including partial or total Internet shutdowns in China, Iran, Egypt, Libya, and Syria; many hundreds of documented DDoS, hacking, and other cyber attacks against political sites; continued growth in the number of countries that filter the Internet; and dozens of well documents cases of on- and offline persecution of online dissidents. The energy dedicated to these battles for control of the Internet on both the government and dissident sides indicated, if nothing else, that both sides think that the Internet is a critical space for political action. In this paper, we offer an overview of our research in the context of these changes in the methods used to control online speech, and some thoughts on the challenges to online speech in the immediate future.

Both sides of the contest, those that seek to control Internet activity—typically but not exclusively governments—and proponents of Internet freedom, have a number of different strategies and tools at their disposal. Repressive governments have a wide range of tools available to them, including technological approaches such as filtering, surveillance and cyber-attacks; information campaigns; and traditional offline methods such as the threat of legal action, physical intimidation and arrest. For advocates of free speech, political and diplomatic efforts to convince governments to allow and protect Internet freedoms are part of a long-term strategy but ultimately rely upon governments to change their policies voluntarily. In the short-term, improving technological tools designed to counter government is among a short list of alternatives and has therefore garnered increasing attention over the past two years. Circumvention tools designed to counter government Internet filters play a prominent role in the clash between censors and freedom of speech advocates. Anonymity tools that aim to counteract surveillance and help to protect user privacy are often lumped together with circumvention tools, although the functional role that they serve is markedly different. Secure hosting to mitigate cyber-attacks is another approach to resisting attempts to limit online speech. Traditional forms of popular advocacy seek to protect activists from arrest and detention and to prevent their disappearance when they are arrested.

One of the puzzles of this field is, given the importance of the Internet as a space for political action, why so few people in filtered countries use circumvention tools to access blocked content. In our 2010 study of the usage of circumvention tools, we found that, at most, only 3% of users in countries with pervasive Internet filtering regularly use circumvention tools.[1] Part of the reason that so few users use the tools is certainly because all of the tools suffer from trade-offs in speed, security, usability,

[1] Hal Roberts et al., "2010 Circumvention Tool Usage Report," Berkman Center for Internet & Society, October 2010, http://cyber.law.harvard.edu/sites/cyber.law.harvard.edu/files/2010_Circumvention_Tool_Usage_Report.pdf.

availability, and accuracy. We evaluated 19 of these tools in 2011 and documented these trade-offs, including that even the fastest of the tools is a few times slower than a direct Internet connection, that many of the tools introduce errors in a significant portion web page requests, and that China especially has become much more successful at blocking many of the tools.

An alternative explanation for relatively low use of circumvention tools is that the tools do not meet one of the major needs of users: creating content on local platforms for local audiences. In our 2011 international survey of politically-oriented bloggers, we found that, for users in filtered countries, the most common reason for not using circumvention tools was simply that they had no need to access filtered content.[2] We also found that a majority of the surveyed bloggers perceived themselves at risk of arrest or persecution for posting political content. Many posted some (but not all) risky content anyway. This finding suggests that projects that focus on providing unfiltered access to international websites are not sufficient. International platforms like YouTube, Twitter, and Facebook that allow local communities to interact and post content certainly play an important role in political activism in repressive countries. But the most important space in this battle may not be in the firewalls between filtering countries and the rest of the world but rather within the local communities in the country where people prefer local over foreign content and daily struggle with decisions about how to self censor to minimize risk of offline persecution.

Our research on Internet locality supports this suspicion. Our finding that roughly 95% of web page requests in China are to sites that are hosted within China bolsters the case that China's most effective form of Internet control has been not only shutting out foreign sites, which it cannot control directly, like Facebook, YouTube, and Blogger, but also fostering the growth of local sites like Baidu, QQ, and Youku that offer the non-political content, community, and functionality that has been the engine of the growth of the Internet everywhere.[3] By spurring the growth of these local sites, China maintains the ability to directly regulate content on the sites while allowing its citizens to access the vast majority of social media content that is not politically controversial.

There are multiple reasons Chinese users choose Youku and QQ over YouTube and Facebook. China's aggressive blockage of these sites is one. The high quality of the Chinese sites and their linguistic accessibility to Chinese users is another. National pride and a desire to use local products may be a third. But the result of these intersecting factors has been the thorough segregation of the Chinese Internet from the rest of the world.

Our 2011 study on the structure of national networks of autonomous systems (the Internet service providers, very large content providers, and other large organizations that route traffic and largely determine policy on the Internet) confirmed the direct, top-down control that China exerts over its

[2] Hal Roberts, et al., "International Bloggers and Internet Control," Berkman Center for Internet & Society, August 2011, http://cyber.law.harvard.edu/publications/2011/International_Bloggers_Internet_Control.

[3] Hal Roberts, "Local Control: About 95% of Chinese Web Traffic is Local," August 15, 2011, https://blogs.law.harvard.edu/hroberts/2011/08/15/local-control-about-95-of-chinese-web-traffic-is-local/.

network.[4] That study found that China uses only four autonomous systems (ASNs) to connect 90% of its 240 million IP addresses (and has only 177 ASNs total) , suggesting that control of a very few "chokepoints" can be used to assert quite thorough Internet control. On the other end of the spectrum, Russia uses 19 ASNs to connect only 30 million IP addresses (and has over 2300 ASNs total), network structure much more complex both absolutely and per capita. The complexity of the network structure may explain why Russia has not filtered its network in the ways China has. But Russia does reportedly engage in a number of other forms of Internet control that reflect this more complex network structure, including DDoS, hacking, and other cyber attacks; on- and offline harassment of activists; and mobilization of youth brigades to flood online forums with pro-government views.

Even though China and Russia differ in which forms of control they use most strongly, a common theme through all of our work has been that China, Russia, and all authoritarian countries we've studies use a diverse set of tactics to control the Internet. Our 2011 report on distributed denial of service (DDoS) attacks against independent media sites found that sites that experienced DDoS attacks usually experienced some other form of Internet control as well, such as filtering, intrusion, or defacement.[5] And the victims of DDoS attacks that we surveyed ranked offline persecution as more serious than online attacks or controls of any sort.

Over all, our work suggests that the increasing complexity of Internet control regimes should force us to rethink our approaches towards empowering Internet users in less open societies. In this paper, we will describe in more detail the studies we have mentioned above and how they fit into this larger story about the diversity of tactics used by autocratic countries to control the Internet. We will conclude by offering some high level recommendations for supporting the work of activists battling these forms of Internet control.

CIRCUMVENTION USAGE

In 2010, we conducted a study to estimate the number of people using circumvention tools worldwide.[6] We used a variety of methods to estimate usage of three different classes of tools: blocking resistant tools, simple web proxies, and virtual private network services. Blocking resistant tools use different methods to evade censorship, but all include sophisticated mechanisms that make it more difficult for a filtering country to block them. This class includes three of the best known circumvention tools: Tor, Ultrasurf, and Freegate. In the simple web proxy class, we included all tools that provide proxied Internet access through a web page interface. Virtual private network (VPN) services provide a virtual

[4] Hal Roberts, et al., "Mapping Local Internet Control," Berkman Center for Internet & Society, 2011, http://cyber.law.harvard.edu/netmaps/mlic_20110513.pdf.

[5] Ethan Zuckerman, et al., "2010 Report on Distributed Denial of Service Attacks," Berkman Center for Internet & Society, December 2010, http://cyber.law.harvard.edu/publications/2010/DDoS_Independent_Media_Human_Rights.

[6] Hal Roberts et al., "2010 Circumvention Tool Usage Report," Berkman Center for Internet & Society, October 2010, http://cyber.law.harvard.edu/sites/cyber.law.harvard.edu/files/2010_Circumvention_Tool_Usage_Report.pdf.

networking device on the client's computer that tunnels all of the user's network traffic through an encrypted tunnel to the VPN proxy. We did not include a fourth class of tools—HTTP/SOCKS proxies—because we have no reliable method to measure their usage. Our survey data and anecdotal reports suggest that use of these tools is low in comparison to other methods.

To estimate the usage of blocking resistant tools, we either surveyed tool developers for their self-reports of usage or, for one tool, used a previously published estimate of usage based on publicly available data. For simple web proxies, we used web page traffic statistics from Google AdPlanner. And for VPN services, we used web searches to compile a list of as many VPN services as we could find and then surveyed all of them for usage data. Of the three classes of tools, we found the most users by far in the simple web proxy class of tools. Through all methods, we estimated an upper range of about 19 million monthly users of the three classes of tools we could measure. Using the OpenNet Intiative's monitoring of national Internet filtering worldwide, we estimated 562 million Internet users in countries with substantial filtering.

If we assume that all circumvention tool users are in filtered countries (which we know that they are not), our usage estimates indicate that at most 3% of users in all filtering countries use one of the three measured classes of circumvention tools. We believe the actual usage number is significantly smaller. Many of these tools are used by students in the US who want to avoid filters on their school's networks and access Facebook, employees circumventing corporate firewalls, or users who want to access content not available in their country, like Hulu's streaming video service. We believe the percentage of Internet users in less open societies using these tools could be as low as 1%.

In addition to collecting data on these individual tools, we looked at Google search frequencies for twenty proxy or circumvention related terms in nine countries in both English and the local language. Through this search term frequency analysis, we found that proxy-related searches are relatively uncommon in all countries, that search for generic versions of the word 'proxy' were vastly more popular that searches for specific tools, and that there were no tools with significantly more search popularity than indicated by their estimated usage numbers.

CIRCUMVENTION TOOL EVALUATION

In February and March of 2011, we conducted an evaluation of nineteen circumvention tools, including simple web proxies, VPN services, and specialized tools, which we identified through a survey of circumvention users (described below). That study is not yet publicly available, as we are redacting it to ensure we do not provide undue assistance to censorious governments. We tested the utility, speed, and accuracy of the tools by using testing servers in China, South Korea, Vietnam, and the United Arab Emirates. For each tool in each country, we requested about 40 sites, half of which were the most popular general sites in the country and the other half of which were sites that were blocked within the country, as determined by the Open Net Initiative or Herdict projects. We used each tool to request each site in each country during each of two separate rounds of tests, one in March and one in February.

Only two of the tools successfully passed both rounds of basic functionality tests in all countries. Two other tools passed one or two tests in a single country but were incompatible with the testing setup and so untestable in other countries. The rest of the tools were either badly broken (meaning they often failed to return a usable web page) or were blocked within at least one test in one country. In the case of VPN tools, this points to the limits of our testing platform, which makes it difficult to fairly evaluate these tools. But in the case of other tools, the results are disturbing. They contrast sharply with the results of a similar evaluation we conducted in 2007, in which we found that the vast majority of tools specifically designed for circumvention were functional and unblocked in all tested countries. We also found that circumvention tools continue to create a significant speed penalty in users' browsing sessions and that many of the tools exhibit high error rates (up to 22%) in fetching and rendering web pages.

These test results are consistent with reports of increasing efforts by some governments over the past two years to block the use of circumvention tools.

CIRCUMVENTION USER SURVEY

In December of 2010, we conducted a survey of politically and internationally oriented bloggers from eighteen countries, including fifteen that substantially filter their Internet connections.[7] The sample of bloggers consisted of blogs that had been linked to by Global Voices Online (GVO), an international community of bloggers who report on citizen media worldwide and that had been tagged by GVO as being associated with one of the eighteen target countries. Of this sample, we found that 57% of respondents regularly used circumvention tools, a much higher number than the 1-3% of general Internet users in filtered countries that we found in the circumvention usage report. This finding was not surprising, however, because the sample for this survey consisted of set of people much more technically skilled and internationally and politically oriented than the average Internet user and, therefore, more likely to use circumvention tools. Perhaps a more surprising finding was that, among the respondents who lived in filtering countries but did not regularly use circumvention tools, the most common reason for not using them was a lack of need to access filtered content. This was by far the most common response to a question about not using circumvention tools, rather than an inability to access these tools, or a lack of knowledge about using them effectively.

The survey also asked a set of questions about users' perceptions of and reactions to the risk of online activism. We found that 74% of respondents perceived some risk of detention, arrest, or criminal investigation in posting material critical of their governments online, and 59% perceived some risk of violence directed at themselves or their families. In response to these risks of posting online, 59% of the respondents had chosen not to post some content online, and overwhelmingly the most common type of content that respondents refrained from posting was political (89% vs. religion at 31% as the next most common). Of those 59% who had censored themselves due to perceived risk, 68% had posted some risky content. We saw no evidence of users who believed themselves to be taking a risk in using

[7] Hal Roberts, et al., "International Bloggers and Internet Control," Berkman Center for Internet & Society, August 2011, http://cyber.law.harvard.edu/publications/2011/International_Bloggers_Internet_Control.

circumvention tools to access content blocked locally. That finding suggests that fear of persecution for using circumvention tools is probably not a factor that explains their low usage.

MAPPING LOCAL INTERNET CONTROL AND LOCALITY

In 2010 and 2011, we analyzed the structure of national networks of autonomous systems (ASNs) to identify the set of ASNs that act as points of control for each national network and to compare the relative complexity of national ASN networks to one another.[8] We defined the points of control for each country to be the smallest set of ASNs that connect 90% of the country's IP addresses. And we defined a measure of complexity that assigned higher complexity to those countries with more ASNs per IP address and those with more IP addresses farther away from the points of control.

We found that in almost every country, only a small subset of ASNs serve as points of control, but that the number of points of control and the relative complexity of national ASN networks differs significantly between countries. Most strikingly, we found that China and Russia structure their network very differently. In our latest analysis, using 2011 data, China had only four points of control for its 240 million IP addresses whereas Russia had nineteen points of control for only 30 million IP addresses. Our measure of complexity likewise indicated vastly different network structures in the two countries, with China scoring as 176 times more complex than Russia, meaning that Russia has an order of magnitude more ASNs per IP address than China and that these IP addresses are much more likely to be located at the edge of its network, away from the core points of control.

These differing network structures reflect the different ways that China and Russia approach control of the Internet. China approaches the problem from a top-down direction, starting with a national program of filtering its Internet connection implemented by a small set of ISPs. This control is complemented by nurturing large-scale national alternatives to blocked international platforms. Russia uses a more decentralized approach to the problem emphasizing end-user surveillance, youth brigades of pro-government commenters, and easily deniable third-party denial of service attacks. Other countries mix top-down and more decentralized strategies to various degrees, including Iran, which uses a combination of pervasive national filtering and frequent third party denial of service attacks to control its network.

We extended this network mapping work to include data about the most popular websites in each country, as reported by Google's AdPlanner service. We found that in both China and Russia, over 95% of web page visits to the most visited 250 sites in the country were to sites hosted within the given country. This is an important finding for understanding Internet control in these countries because those locally hosted sites are subject to traditional, knock-on-door regulation. To control the content on a site like Baidu.com, China needs only send a government agent to the door of the person who runs the

[8] Hal Roberts, et al., "Mapping Local Internet Control," Berkman Center for Internet & Society, 2011, http://cyber.law.harvard.edu/netmaps/mlic_20110513.pdf.

site and threaten to fine, arrest, or otherwise enforce its policies, while China's ability to control what appears on Facebook.com is significantly more limited.

China's blockage of many of the big international platform sites can explain some of this high locality of web traffic, as can isolation through language, or cultural preferences for local services. Whatever the causal reason, the end effect is the same: a preference for locally hosted content makes local knock-on-door enforcement regulation highly effective at controlling the Internet. The Russian case is even more interesting because Russia does not filter any of the big websites and does not exert control as directly over its local content providers. However, the locality of its web traffic through cultural or linguistic factors, makes its content publishers vulnerable to local persecution, for example, by harassment of sites under a law that makes publishers responsible for inflammatory content posted in discussion forums.

DDOS ATTACKS AGAINST INDEPENDENT MEDIA

Our 2010 research on distributed denial of service (DDoS) attacks against independent media used four methods to explore the prevalence, form, and recommended response to these attacks: we conducted an extensive media review for reports of politically motivated DDoS attacks; we surveyed 317 independent media organizations in nine countries; we followed up the survey with detailed technical interviews of operators of twelve independent media sites that had suffered from DDoS attacks; and we held a working meeting of technologists, independent media publishers, academics, and human rights organizations.[9]

The themes that ran through the results of all methods employed were that DDoS attacks are common against independent media sites, that most independent media sites suffer from an array of different controls in addition to DDoS attacks, and that the best defense for sites at high risk for these attacks is to seek protection from one of the few dozen giant companies that host services that are core to the modern Internet, like Google, Facebook, Amazon, Akamai and a few others.

The media research found reports of 140 politically motivated attacks against 280 sites, and we think that this strongly under-reports the number of attacks taking place because we focused on English language media and because most attacks are not reported in the media. Our survey found that, among our respondents, 81% of sites that had suffered from DDoS attacks had also suffered from at least one incidence of filtering, intrusion, or defacement. And respondents to the interview ranked filtering and offline persecution to be bigger problems than DDoS attacks. Follow-up technical interviews reinforced this finding that most sites experience of range of different on- and offline controls and that DDoS attacks (and filtering) are often not the most important types of controls. From our working meeting, we learned that defense against DDoS attacks and other kinds of Internet controls is often very difficult for an independent organization, and that a key part of the solution to these attacks is connecting local

[9] Ethan Zuckerman, et al., "2010 Report on Distributed Denial of Service Attacks," Berkman Center for Internet & Society, December 2010, http://cyber.law.harvard.edu/publications/2010/DDoS_Independent_Media_Human_Rights.

communities connected to independent media to the core Internet companies that have the technical resources to help them.

CONCLUSIONS AND RESPONSES

Online speech in less open societies faces a very different climate now than it did just three or four years ago. While national Internet filtering of web traffic has become more pervasive, this tactic may be less relevant than aggressive control of locally published online content. Filtering is no longer the only technique employed by opponents of online speech. Publishers need to worry about DDoS attacks, site hijacking and defacement, punitive enforcement of defamation and libel statutes, including intermediary liability for comments, as well as laws that prohibit dissemination of "national secrets". Governments are taking unprecedented steps to control online speech, including intercepting passwords to social network services, as in Tunisia, and shutting down the entire Internet, as in Libya and Egypt.

Our research suggests that Internet users and publishers of online content are often ill-prepared to cope with these shifts in Internet control. A small fraction of individuals who experience Internet filtering are using tools to circumvent censorship, and those tools may be less effective than in years past. And while new tools have emerged to help users evade censorship, there's little hope that a technical "fix" will solve problems like domain name hijacking or DDoS.

Our thinking about technological solutions to internet control has changed sharply in the face of this changed environment. Four years ago, we were reasonably sure that the developers of circumvention tools were winning the match against government censors. Not only is victory in that match less assured now, the entire playing field has changed, and new technologies of control are far harder to defend against than national internet filtering.

While circumvention and anonymity tools are likely to continue to play an important role in digital activism in repressive online environments, investing disproportionate resources on technological solutions may be counter-productive. Increasing the performance and availability of circumvention tools may make it easier for a larger set of Internet users to access content that would otherwise be out of their reach, just as reducing the effective cost of using anonymity tools could make it safer for political dissidents to express themselves online and investments in more secure hosting solutions would reduce the vulnerability of independent media sites to malicious attack. It is unclear though how advances in each of these 'liberation technologies' might contribute to freedom of expression online or political reform, particularly given that none of these changes would take place in a vacuum; to the extent that they succeed, all of these actions will shift the benefits and costs to repressive governments of implementing stricter Internet controls and investing in more sophisticated counter-measures. Defeating government Internet filters, even if feasible, may provide the catalyst for more draconian government restrictions. In such a context, it is impossible to predict whether these moves will ultimately increase or decrease online access to information and the ability of people to organize online.

In the wake of these changes, we offer five suggestions for everyone engaged in the struggle for an open internet.

FOCUS ON CIRCUMVENTION TOOLS FOR ACTIVISTS. Our research suggests that it's unrealistic to believe that circumvention tools will be used by a very broad audience, given a preference for local content in some markets and the ongoing challenges in making circumvention tools fast and easy to use. We believe that focusing on building highly reliable, blocking resistant tools with fast throughput for a small audience of users might focus the attentions of software developers more precisely. We believe there's some truth to Xiao Qiang's idea that a small set of internationally connected activists can disseminate information through local networks. And we worry that efforts to reach very broad audiences with circumvention tools is trying to solve a demand problem by focusing on supply.

ADDRESS DDOS AND OTHER CHALLENGES. Problems like DDoS seem like ones that affect publishers, not internet users. But that's the subtlety of this form of control. By disabling a site, hijacking a domain name or otherwise disabling a web presence, governments make content inaccessible to users around the world, not just users in their own countries. The tools and systems we have to respond to DDoS, site hijacking and other technical threats are poor, and could be vastly improved. The solution is probably not a purely technical one – it involves building teams of experts who can support publishers in less open societies with best practices, incident response and access to less vulnerable, shared platforms. There's a great deal of work to be done in this space, and a few promising efforts that could benefit from support.

UNDERSTAND THE TECHNIQUES NEEDED TO WIN LOCAL BATTLES. Many of the ways in which the Internet is controlled are highly local. While we like to pretend that the network is global and seamless, it's very important to acknowledge that Chinese users use Chinese sites… and that those sites are tightly controlled by forces within China. When Internet traffic is local and threats to publishers are local, we may need fewer technical responses to Internet control, and more responses analogous to those used by press freedom organizations: naming and shaming, censorship indices, pressure through international bodies, and campaigns to protect individual dissidents. While technical solutions that promise liberation from these local controls are very compelling conceptually, we believe many of the necessary responses are the sorts of messy, ground-level human rights work that's quite unfamiliar to the technical community.

SEEK NEW ALLIES. The tools built to allow millions of users to circumvent Internet censorship have been built by small, poorly-resourced teams around the world. With the adoption of an "Internet Freedom agenda" by the US State Department, a new pool of money has appeared to support this work. But the most powerful potential allies in a battle for an open internet are still mostly on the sidelines. Companies like Google, Facebook, Amazon, Akamai, Microsoft and others have resources that could be marshaled to the benefit of Internet users and publishers in closed societies: bandwidth capacity that could support circumvention systems, DDoS-resistant hosting that could protect publishers; technical expertise that could fend of domain hijacking and intrusion. Helping these companies come off the sidelines and bringing them into the game is a key challenge for the future of free speech online.

MONITOR IN REAL TIME. Because the conditions for online speech are changing so rapidly, assessment of speech environments on an annual basis is increasingly insufficient to understand the challenges at hand. Our recent round of testing of circumvention tools hints at the disturbing possibility that some governments are periodically attempting to block traffic based on protocol and traffic pattern, potentially disrupting entire classes of circumvention tools. We need to develop a strategy for monitoring filtering, tracking DDoS and intrusion attacks and documenting assertions of local control in a way that is timely and ongoing to fully understand the threats to an online speech environment. Implementing this type of project might be a first step that involved cooperation between corporate actors, traditional free speech advocacy groups and the technology experts who have been tracking threats to free expression online.