



Internet Censorship and Control

Steven J. Murdoch
University of Cambridge

Hal Roberts
Harvard University

The Internet is and has always been a space where participants battle for control. The two core protocols that define the Internet – TCP and IP – are both designed to let separate networks connect to each other easily, so that networks that differ not only in hardware implementation (wired versus satellite versus radio) but also in their politics of control (consumer versus research versus military) can interoperate easily. It's an Internet feature, not a bug, that China – with its extensive, explicit censorship infrastructure – can interact with the rest of the Internet. However, where there is interaction crossing social and political boundaries, there will be conflict as to what controls should be in place, and who should enforce them.

Who Governs the Internet ... and How?

Initially, Internet control mechanisms were developed largely by members of the Internet community in the form of norms enforced with little or no intervention from courts or law enforcement. Community members, especially technical administrators, enforced these

norms by threatening temporary or permanent disconnection. For example, the “Usenet Death Penalty” (see <http://catb.org/jargon/html/U/Usenet-Death-Penalty.html>) was a punishment that could be imposed on ISPs, preventing their users from posting to Usenet, if the consensus among major ISPs' technical administrators was that the offender wasn't adequately controlling spam.

Many hoped that these internally developed and enforced methods of controlling Internet users would be sufficient to preserve the Internet's ability to function well. By avoiding government interference, the ambition was to achieve freedom of speech and freedom from prejudice. These goals were represented in John Perry Barlow's “Declaration of the Independence of Cyberspace,” in which he proposed dealing with the conflicts that existed at the time through a shared, international “social contract,” rather than by extending national laws to apply to the Internet (see <https://projects.eff.org/~barlow/Declaration-Final.html>).

Increasingly, however, governments have imposed Internet control

mechanisms – including technical, legal, political, and social tools – due to a perception that self-regulation is no longer sufficient to deal with challenges resulting from the rapidly growing number and diversity of users, intensifying criminal activity, the Internet's increasing role as a core social infrastructure, and the diversity of participating countries' political philosophies.

M. Christopher Riley's "Anarchy, State, or Utopia? Checks and Balances in Internet Governance" gives an overview of Internet control mechanisms to give context to the discussion of censorship and control in the other articles presented. Riley's article shows how a complex set of checks and balances, spread among governments, international bodies, companies, and individuals, has evolved over time and is continuously in flux.

Riley describes how the idealistic vision of Internet self-governance failed to sufficiently deal with more modern challenges, leading other bodies to step in to try to deal with the problems faced. Such efforts haven't been easy, though; no universal agreement exists on what Internet governance goals are, let alone how to achieve them. One particular concern is that this evolution of control mechanisms will reduce the Internet's power as a force for freedom and turn it into a net force for government repression.

Technological Solutions and Pitfalls

There are reasons for believing that this dystopia is possible – enforcement of Internet controls can be automated, allowing authorities to exercise such controls at far larger scales than possible with other forms of communication. A government that can monitor an entire country's population is a nightmare for civil libertarians and a dream of dictators. Achieving this for the Internet could be as easy as enabling the right configuration options on existing networking equipment. The aftermath of the Arab Spring revealed how governments targeted their surveillance at civil society, and how Western companies were knowingly complicit in supplying these governments with the technology to do so.

Encryption and authentication can help Internet users resist surveillance, but these techniques are underused, hampering the potential for security improvements. Some of this failing

is due to governments classifying encryption systems as munitions and restricting their distribution. Some is a result of the difficulty in correctly using encryption software. Consequently, people who do use it are more likely to stand out in the crowd and so put themselves at risk for more targeted surveillance that even advanced encryption systems can't withstand. Again, Western companies have been selling to authoritarian governments software packages that circumvent encryption by exploiting vulnerabilities in commonly used software.

One area in which encryption has been widely adopted, however, is HTTPS-encrypted Web browsing. Initially motivated by the goal of achieving safe Internet commerce, HTTPS is now applied to encrypt access to webmail services. Webmail over HTTPS doesn't offer the same level of security as end-to-end email encryption such as OpenPGP. HTTPS encrypts only the communication between the user's Web browser and the webmail provider's servers, not the traffic between mail servers. But HTTPS is easier to use than more secure email encryption technologies. In fact, it requires no extra software, and most users won't even realize the difference between HTTPS and plain HTTP webmail. Moreover, unlike OpenPGP, users can adopt webmail over HTTPS unilaterally, without having to persuade all of their communication partners to upgrade too.

The trend among webmail providers to offer HTTPS-encrypted access to their services is encouraging because it can help users avoid government monitoring (though it won't protect the users from the webmail provider itself). However, the control mechanisms that were put in place to protect HTTPS users are increasingly inadequate. Certification authorities (CAs) are responsible for establishing that the HTTPS site a user is connecting to is, in fact, run by the owner of the site's domain name. CAs effectively have the power – with their cryptographic credentials baked into Web browser code – to impersonate any website to any user. This power has made CAs targets for attacks from criminals, who have successfully obtained fraudulent certificates in a few cases. The possibility that governments will pressure CAs into granting fraudulent certificates to disguise surveillance operations is also a concern.

Steven Roosa and Stephen Schultze explore these topics in "Trust Darknet: Control and

Compromise in the Internet's Certificate Authority Model." The authors have examined a spate of compromised, high-profile CAs and illuminate weaknesses in the CA model that have made such compromises so damaging and hard to manage. Decentralizing CAs has encouraged competition, which has pushed down prices but led to a race to the bottom in terms of security. When coupled with technical limitations that mean a compromise of any CA leads to every website being at risk, the situation is far from ideal. The authors examine the legal and economic forces at work and discuss improvements that they hope will reduce the likelihood of CA compromise and the damage that would result if one does occur.

Measures under consideration include forcing CAs to be more transparent about how they grant certificates and how browser vendors

restricted to the most repressive countries, were easily defeated. But governments learned from their mistakes, and today's censorship techniques are proving increasingly effective and are gaining widespread use – that is, they're adopted by dictatorships and democracies alike. Techniques for censorship vary, ranging from directly interfering with Internet traffic to pressuring content providers to remove offending material. Various motivations for censorship arise, such as political control, child protection, and revenue protection for copyright holders. Frequently, however, a censorship system introduced for one reason might later be used for another, as when the UK's system for blocking images of child sexual abuse was used to block the Pirate Bay BitTorrent search engine.

In "Censorship v3.1," Derek Bambauer examines how censorship's nature has changed over time. While this change has been taking place, international bodies have been fighting for more control over the Internet, seeking to reduce the US's influence over key Internet decision-making bodies. Bambauer's article discusses approaches to reducing the harm that can come from Internet censorship – including recognizing that restricting access to material is censorship (although perhaps defensible) – supporting Internet governance's decentralized nature, and resisting efforts to outsource censorship to companies that are less accountable than governments – all of which aim to get discussions around censorship out in the open.

Once society admits that censorship is taking place, the debate can move on to topics such as whether the proposed censorship is proportionate, who has jurisdiction when standards vary between countries, and what checks and balances should be put in place. Achieving transparency for Internet censorship has proven challenging, even in mature democracies. Reasons given for not disclosing lists of blocked sites include claims that listing them as censored will increase attention paid to them and that transparency isn't required when censorship – rather than coming directly from the government – is outsourced through informally imposing pressure on content hosting companies or access providers.

For these reasons, researchers are largely left the task of revealing the extent of censorship; such researchers seek to discover which content is being blocked, by whom, and for what reasons.

Sometimes, over-blocking is an attempt to avoid criticism, but other times it proves to be a mistake resulting from overzealous interpretations of rules.

include CA credentials in Web browsers. In this way, the Internet community can detect a rogue or compromised CA that issues a certificate to the wrong person, though perhaps only after the damage has been done. Detecting misbehavior can at least get the CA blacklisted in browsers and prevent future mistakes. Being blacklisted by a commonly used browser will destroy a CA's business model, so significant incentive exists for avoiding this fate. These measures' effectiveness remains to be seen, but we can learn from the experience of trying to introduce transparency to another type of control on the Internet – censorship.

Making Censorship Transparent

Initially, censorship on the Internet was viewed as futile, with Internet pioneer John Gilmore famously saying, "The Net interprets censorship as damage and routes around it."¹ Governments' first crude attempts at censorship,

Although technical studies play an important role here, their results aren't themselves adequate to give a full picture. Knowing which sites are blocked or which content has been removed is necessary, but can't always reveal why censorship has occurred. Often, sites are blocked even though they fall outside the stated criteria for censorship. Sometimes, this over-blocking is an underhanded attempt to avoid criticism, but other times it proves to be a mistake resulting from overzealous interpretations of rules or collateral damage due to technical limitations in censorship techniques. Distinguishing these cases is important for informing the debate, because if errors happen too frequently, then arguments for proportionality in censorship could be less valid.

"Not By Technical Means Alone: The Multi-disciplinary Challenge of Studying Information Controls," by Masashi Crete-Nishihata, Ronald Deibert, and Adam Senft, describes the mixed-method approach that the OpenNet Initiative took for building a global survey of Internet censorship. This study combined technical measurements with analysis of the political, legal, and economic systems behind information controls. The article discusses the technical and methodological challenges of conducting the study and illustrates the OpenNet approach through various case studies.

Casting light on censorship can be particularly challenging when governments deliberately try to disguise the type of censorship they perform. One way this occurs is when governments don't directly censor content but use laws and intimidation to cause individuals to self-censor. Surveillance, or at least the perception thereof, gives its targets a realistic expectation that if they step out of line, they will be at risk. Frequently, surveillance is coupled with some content removal or blocking so that people are aware that their online activities are being monitored, and the limits of what's considered acceptable are clear. The technical measures set the limits, and the risk of punishment keeps individuals from testing them. As such, just because something isn't blocked doesn't mean that many people will exploit this freedom.

In "Assessing Censorship on Microblogs in China: Discriminatory Keyword Analysis and the Real-Name Registration Policy," King-wa Fu, C.H. Chan, and Michael Chau examine self-censorship. Through monitoring the

censorship of posts on a popular Chinese microblogging service, they infer which topics censors consider sensitive. From this, they discuss how the use of sensitive terms might have changed when China introduced its real-name registration policy, which increased accountability for microbloggers and might have created a chilling effect among those discussing controversial topics.

The articles presented in this issue make it clear that no global consensus exists on what mechanisms of control are best suited for managing conflicts on the Internet, just as there is none for other fields of human endeavor. That said, we can be optimistic that with vigilance and continuing efforts to maintain transparency, the Internet can remain a force for increasing freedom rather than become a tool for more efficient repression. □

Reference

1. P. Elmer-Dewitt, "First Nation in Cyberspace, *TIME Int'l*, 6 Dec. 1993; www.chemie.fu-berlin.de/outerspace/internet-article.html.

Steven J. Murdoch is a Royal Society University Research Fellow at the University of Cambridge Computer Laboratory and a fellow of Christ's College. He conducts research on developing metrics for computer security and on how to design safe communication systems for people living and working in repressive regimes. His research also includes developing technology that circumvents censorship, and improving bank payment system security. Murdoch has a PhD in computer security from the University of Cambridge. Contact him at steven.murdoch@cl.cam.ac.uk.

Hal Roberts is a fellow at the Berkman Center for Internet & Society at Harvard University, where he is a cofounder of the Media Cloud project, an open platform for the collection and analysis of the content of online media. He studies issues around the flow and control of information online. Roberts has written papers on Internet filtering circumvention, Internet surveillance, distributed denial-of-service attacks on independent media, and online media. Contact him at hroberts@cyber.law.harvard.edu.



Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.