

Introduction

Jonathan Zittrain and John Palfrey

A Tale of Two Internets

Tens of thousands of international travelers descended upon the Tunis airport for the World Summit on the Information Society in 2005. The summit brought together policy-makers, journalists, nongovernmental organization (NGO) leaders, academics, and others to consider the present and future of information and communications technologies. Polite Tunisian handlers in crisp, colorful uniforms guided arriving summit attendees to buses that took those with credentials to one of several sites nearby.

The capital, Tunis, hosted the main conference facilities. The seaside town of Yasmine-Hammamet, with boardwalks, theme parks, casinos, and breathtaking sunsets, housed delegates who could not find lodging in the city. Within the main conference facilities in Tunis, they would experience the Internet as though they were in a Silicon Valley start-up: unfettered access to whatever they sought to view or write online.

But those by the sea in Yasmine-Hammamet, outside the United Nations–sponsored conference facilities, encountered a radically different Internet—the one that is commonplace for Tunisians. If attendees sought to view a site critical of the summit’s proceedings or mentioning human rights—for instance, a site called Citizen’s Summit, at www.citizens-summit.org/—they would see a page indicating that a network error had occurred. Among other curious things, the page was written in French, not the native Arabic. The blockpage is partially accurate: something in the network had caused that information never to reach the surfer’s laptop.¹ But it was not an error.

The blockage is intentional, one of thousands put in place daily by the government of Tunisia. The ad hoc filtering of information underway in Tunisia is flatly at odds with the ideals touted by World Summit participants. Tunisia’s filtering system was implemented long before the World Summit kicked off, and it was unaffected by the attention the summit brought to Tunisia.

A filtering system is meant to stop ordinary citizens from accessing some parts of the Internet deemed by the state to be too sensitive, for one reason or another. The information blocked ranges from politics to sexuality to culture to religion. As user-generated content has

gained in popularity and new tools have made it easier to create and distribute it, filtering regimes have pivoted to stop citizens from publishing undesirable thoughts, images, and sounds, whether for a local or an international audience. The system that facilitates a state's Internet filtering can also be configured to enable the state to track citizens' Web surfing or to listen in on their conversations, whether lawful or unlawful.

A Tale of Many Internets

Tunisia is not a special case. More than three dozen states around the world now filter the Internet. This book contains the results of the first systematic, academically rigorous global study of all known state-mandated Internet filtering practices. Previously, the OpenNet Initiative and others have reported only anecdotally or sporadically on the scope of Internet filtering. Our first goal in writing this book is to present the data from this global study, allowing others to make use of it in their own empirical work, or to place it within a normative framework. Second, in addition to state-by-state test results, we have commissioned a series of essays analyzing these test results and related findings from a variety of perspectives—what this emerging story means from the standpoint of technology, as a matter of international law, in the context of corporate ethics, and for the vibrant activist and political communities that increasingly rely upon Internet technologies as a productivity enhancer and essential communications tool.

For this first global study, we have sought to find those places in the world that practice state-mandated technical filtering. The definition of what we are and are not covering here is important to set forth at the outset: we seek to describe technical blockages of the free flow of information across the Internet that states put in place or require others to institute. To determine where to test for such blockages, we have drawn upon our own technical probes and forensic analyses of networks, published reports of others who track these matters, and credible unpublished reports that we received either through interviews or over the transom. Our emphasis on state-mandated technical filtering underscores our own sense that “West Coast Code,” in Lawrence Lessig's terms (computer code), is more malleable, more subtle, more effective in many contexts, and less easily noted, changed, or challenged than “East Coast Code” (ordinary law and regulation), which is typically less opaque in its operation.² Straight-forward state regulation of speech without technological components can, of course, result in censorship; our work here is designed to focus on regulation that, when implemented through code, seems more a force of nature than an exercise of political or physical power.

Thus it is entirely possible that a state that does not require or inspire technical filtering can possess a set of regulations or social norms or market factors that render its information environment less free than a state with fairly extensive technical filtering. A rich and comprehensive picture of what a truly “free” or “open” information environment looks like can rely only in part on conclusions about Internet filtering. The essays that accompany our presentation

of the data are intended to provide some, though by no means all, of the relevant context. A shrewd observer might well make a case that the extensive regulatory regimes for speech in Canada, the United States, and the United Kingdom—from which states the majority of our researchers hail—result in a more constrained information environment than a state with technical filtering but little else by way of law, norms, or markets to constrain an Internet user. We map out filtering practices, and the law and regulation behind them, so that they may take their place within a larger mosaic of assessing and judging the flow of information within and across the world's jurisdictions.

The states that practice state-mandated filtering are predominantly clustered in three regions of the world: east Asia, the Middle East and North Africa, and central Asia. A handful of states outside these regions also encourage or mandate certain forms of filtering. Someone in the United States, for instance, may encounter state-mandated Internet filtering on some computers in libraries and schools. A citizen in northern Europe might find child pornography blocked online. In France and Germany, content that includes imagery related to Nazism or Holocaust denial is blocked in various ways and at various levels. The emerging trend points to more filtering in more places, using more sophisticated techniques over time. This trend runs parallel to the trajectory of more people in more places using the Internet for more important functions in their lives.

We find that filtering implementations, and their respective scopes and levels of effectiveness, vary widely among the states that filter. China institutes by far the most extensive filtering regime in the world, with blocking occurring at multiple levels of the network and spanning a wide range of topics. Singapore, by contrast, despite a widely publicized filtering program, in fact blocks access to only a handful of sites. Each of the sites blocked in Singapore is pornographic in nature. Several states, including some in central Asia, filter only temporarily when elections or other key moments make the control of the information environment most important to the state. Most states implement filtering regimes that fall between the poles of China and Singapore, with significant variation from one to the next. Each of these state-mandated filtering regimes can be understood only in the political, legal, religious, and social context in which they arise. It is just this context that we seek to provide in the chapters that follow our presentation of the data.

Our aim in this volume is to document, with the greatest degree of precision possible, technical Internet filtering wherever we have been able to find it, and to set it in a context that acknowledges the nuances and complexity of this matter. We have relied upon an extensive network of researchers in each of the regions of the world that we have studied, as well as area-studies experts based outside those regions. We chose to study and report on the states covered in this volume, as well as other states that appear not to be filtering but are on our “watch list,” because our researchers, members of the press, or others in this field—Reporters Sans Frontières or Human Rights Watch, for instance—have identified these states as potentially carrying out state-level filtering. The lists used in the testing that forms the core

of our set of findings are the product of study of the political, social, cultural, and religious issues in each of the states we have reviewed. While there is no doubt filtering underway in places around the world that we have yet to uncover, our goal in this volume is to be as comprehensive as possible.

The core of the data we present is found in short reports covering each state that we have studied in depth, with an overview for each of the three regions—east Asia, the Middle East and North Africa, and central Asia—identifying themes and trends across states. The section on testing results for each state sets forth the types of content blocked by category and includes documentation of the most noteworthy content-specific findings.

We intend to update this study annually. Our intention is to develop a publicly accessible online database of filtering test results worldwide over time. Taken together, these reports represent a starting point in understanding the nature and future of global Internet filtering.

In addition to the state-specific data, we present a series of chapters that builds arguments grounded in our empirical findings about Internet filtering. The first short chapter, by Robert Faris and Nart Villeneuve, includes a set of issues that emerge from the data: trends and themes from a global perspective. Our chapter 2 gives an overview of the politics and practice of Internet filtering. The third chapter, by Ross Anderson and Steven Murdoch, considers the technology that powers the Internet filtering and highlights its strengths and limitations. The fourth chapter, by Mary Rundle and Malcolm Birdling, takes up the extent to which international law might bear on Internet filtering. Our chapter 5 examines the ethical issues for corporations seeking to avail themselves of markets in states that filter. The final chapter, by Ronald Deibert and Rafal Rohozinski, looks in depth at the impact of Internet filtering upon the activist community that increasingly relies upon the Internet for mission-critical activities.

While we bring our own normative commitments to this work—those of us who have contributed to this work tend to favor the free flow of bits as opposed to proprietary control of information, whether by states or companies or both—our goal is not to point fingers or assign blame, but rather to document a trend that we believe to be accelerating and to set that trend in context. We seek to prompt further conversation across cultures and disciplines about what changes in Internet filtering practices mean for the future of the Internet as well as the future of markets, social norms, and modes of governance around the world. We look forward to the conversations as others put these data into the proper, broader context—into the larger mosaic of political and cultural freedom—into which they belong.

Notes

1. For one of many contemporaneous accounts, see John Palfrey, *On Being Filtered in Tunisia, or, What WSIS Should Really Focus On*, <http://blogs.law.harvard.edu/palfrey/2005/11/14/on-being-filtered-in-tunisia-or-what-wsis-should-really-focus-on/>.
2. Lawrence Lessig, *Code and Other Laws of Cyberspace* (New York: Basic Books, 1999), 53–54.