Weaving the Authoritarian Web:

The Control of Internet Use in Nondemocratic Regimes

Taylor C. Boas

Department of Political Science

University of California, Berkeley

tboas@uclink.berkeley.edu

In John Zysman and Abraham Newman, eds., *How Revolutionary Was the Digital Revolution? National Responses, Market Transitions, and Global Technology.* Stanford, CA: Stanford Business Books

In the preparatory meetings leading up to the December 2003 World Summit on the Information Society in Geneva, the delegations of several authoritarian regimes reacted strongly against the hands-off approach to Internet regulation promoted by the United States and other advanced democracies. Saudi Arabia, for instance, proposed that the development of the information society "shall be done without any prejudice whatsoever to the moral, social, and religious values of all societies"—values to which the Saudi government has appealed when justifying its own regime for Internet censorship. The Chinese delegation campaigned strongly against a statement of support for the principals of free speech enshrined in the Universal Declaration of Human Rights. Ultimately, the Summit's final declaration disregarded the objections that these and other authoritarian governments had voiced during negotiations, but their positions stand as a strong statement that not all countries accept a laissez-faire vision for the future of the Internet.

At first glance, the negotiating positions taken by China and Saudi Arabia might suggest that authoritarian leaders in the information age face a stark choice: promote the development of an Internet that remains free from extensive government control, or exert control over the technology by restricting its diffusion. Whether because of inherent technological characteristics that complicate efforts to censor the Internet, or because countries are under pressure to align their policies with those preferred by the international community, many scholars have assumed that the only effective way to control the Internet is to limit its growth or even keep it out entirely. Milner (2003a, 2003b), for instance, hypothesizes that authoritarian leaders will be less likely than democratic ones to promote Internet development, and she uses indicators of diffusion (such as users or hosts per capita) as proxies for government policy toward the Internet. Franda (2002) interprets national policies to restrict the free flow of information as being deviant and “isolationist” with respect to the international regime for Internet governance, and he expresses skepticism that they will be sustainable. Kedzie (1997) argues that the technology poses a “dictator’s dilemma” to autocrats who must either connect to the Internet and democratize, or shun the information revolution and accept economic decline.

While its use can undoubtedly pose challenges to authoritarian rule, the Internet is an attractive technology to all governments, democratic and authoritarian alike, and hardly any dictator has been willing to ignore it entirely. Internet diffusion offers substantial economic benefits in terms of the potential development of an e-commerce sector, establishing conditions conducive to foreign investment, and stimulating existing domestic industries. Use of the Internet within government itself (e.g., procurement) can facilitate the development of a rational bureaucracy, reducing opportunities for corruption and graft. Implementing government services online (payment of income taxes and the like) can increase efficiency and boost public satisfaction with the regime. If it is possible for authoritarian rulers to have the best of both worlds—reaping the benefits of Internet diffusion while staving off any potentially destabilizing political effects—they will certainly want to do so.

In this paper I argue that, contrary to the assumptions of many studies of the Internet in authoritarian regimes, governments can establish effective control over the Internet while simultaneously promoting its development. Indeed, China and Saudi Arabia are two of the most prominent examples of this phenomenon. While they may have been less influential than they had hoped in negotiations over the global governance of the Internet, both have long sought to implement within their own borders the principles they recently espoused in Geneva. Far from trying to regulate the Internet by merely restricting its diffusion, these countries have employed both technological and institutional means to control use of the Internet while also encouraging its growth. In doing so, they stand as counter-evidence to much of the optimistic thinking about the Internet and democratization that was voiced by pundits and politicians during the early days of the Internet and the technology boom of the late 1990s.

In the first section of this chapter I address the technological bases of Internet control in authoritarian regimes. Much of the early scholarship on the feasibility of government regulation of the Internet pointed out that the network was initially designed as a technology that would be difficult to control at a centralized level (Froomkin 1997, Johnson and Post 1996; for a review, see Boas 2004). I argue, however, that this control-frustrating characteristic of the early Internet is not necessarily locked into place as the technology diffuses around the globe. On the contrary, the logic of Internet diffusion means that the global network is quite flexible and capable of being modified in new environments, allowing authoritarian regimes to embed control-facilitating technological features into the portions of the global Internet that fall within their borders.

While most authoritarian regimes have exploited the flexibility of Internet technology to implement technological measures of control, determined users have almost always found ways to circumvent these barriers. In the second section of the chapter, therefore, I distinguish between perfect and effective control—the former being what matters for tech-savvy individuals that want to gain unfettered access to the Internet; the latter being what authoritarian regimes actually pursue. It is in establishing effective control over Internet use that institutional constraints on behavior—law, social norms, and the market—come most clearly into play. By manipulating the architecture of a flexible technology, and by leveraging influence over laws, social norms, and the market in ways that supplement these architectural constraints, the leaders of authoritarian regimes can exert control over the use of a supposedly control-frustrating technology.

Throughout the paper I illustrate these conceptual and theoretical arguments about the Internet in authoritarian regimes with evidence from the cases of China and Saudi Arabia. As the two countries that have developed what are probably the world's most extensive technological mechanisms for Internet censorship, China and Saudi Arabia are not intended to be representative of authoritarian regimes as a whole. Rather than showing what is *typical* of non-democratic governments, these extreme cases of Internet regulation illustrate what is *possible*. If each has largely succeeded in establishing control over the Internet, others may prove similarly capable in the future.[1]

Institutional and Technological Constraints on Internet Use

In evaluating the potential for establishing control of the Internet in authoritarian regimes, it is useful to consider the means by which authorities might seek to do so. In his study of Internet regulation in advanced democracies, Lessig (1999) has identified four specific mechanisms—law, social norms, the market, and architecture—that governments can employ to control Internet use. The first three can be loosely grouped together as *institutional* constraints—"the humanly devised constraints that shape human interaction" (North 1990: 3). The manner in which they influence behavior is fairly straightforward: laws threaten punishment for prohibited activities, violators of social norms may incur ostracism, and the market can encourage or discourage particular activities based on their cost. As societal constructs, each of these institutional constraints is capable of evolution and change over time. Laws are challenged and overturned; social norms evolve; markets fluctuate, and the degree to which any individual is constrained by them varies with wealth.

Architectural means of regulation occupy a somewhat different category than institutional constraints. In the case of the Internet, architectural constraints consist of the technological characteristics that make certain types of Internet use easier, more difficult, or impossible. In contrast to institutional constraints on Internet use, the technological architecture of the Internet is not as obviously capable of significant evolution. The Internet is a technology whose diffusion is characterized by increasing returns to scale; historically, many such technologies have been examples of path-dependent development and the lock-in of technological characteristics that remain static over time (Arthur 1994, David 1985). If we accept that the Internet's founding characteristics initially made it difficult to control, and if the diffusion of the Internet does indeed give rise to technological lock-in, then the lack of an effective architectural constraint on Internet use might actually be quite *incapable* of change over time.

If true, the potential persistence of a control-frustrating Internet architecture bears special significance for the regulation of Internet use in authoritarian regimes. When effectively implemented, architectural constraints are the only type of regulation that can exert immediate and absolute control over human behavior (Lessig 1999). Laws and social norms can be violated at will; sanctions for such violations are imposed by a government or community only after the fact. Market constraints can be violated in the form of theft; market actors must rely on both social norms and the legal system for effective enforcement. But a technological architecture that makes certain types of Internet use impossible cannot be circumvented even at the risk of future sanctions, and the effectiveness of this constraint does not depend on support from the community or the legal system.

Conversely, if the Internet's architecture is inherently unable to prevent certain types of online behavior, it is impossible for governments to place absolute constraints on Internet use.

The combination of law, social norms, and market constraints can discourage the prohibited activity, but they can never render it impossible. Thus, the supposed rigidity of the Internet's technological architecture is a cornerstone of the argument that the medium inherently frustrates governments' efforts at control. To determine whether the development of the Internet in authoritarian regimes does in fact involve the replication of its initial control-frustrating characteristics, it is useful to see how well the dynamics of path dependence describe this technology's global diffusion.

Path Dependence and the Internet's Control-Frustrating Characteristics

The concept of path dependence in technological development describes a pattern in which the particular configuration for a new technology becomes "locked in" over time as increasingly widespread use raises the cost of switching to another alternative. In particular, the diffusion of such technologies involves increasing returns to scale, which derive from at least one of several characteristics (Arthur 1994). The technology may have a large ratio of fixed to marginal cost, so that the production cost per unit declines as production increases. The technology's adoption may also be characterized by learning effects—the more it is used, the more its efficiency can be improved vis-à-vis other alternatives. Finally, path dependent technologies often display network effects, in which the demand for the technology (and its value to each current user) increases with each additional unit sold.[2] The chosen technology constitutes a standard around which users coordinate, and while any one of them might *ceteris paribus* prefer a different technological configuration, the benefits of standardization outweigh the benefits of switching.[3]

The Internet shares each of these characteristics, making it a technology whose adoption generates increasing returns to scale. Establishment of the Internet's physical infrastructure and development of its core protocols involved significant fixed costs, which were underwritten by both the U.S. government and AT&T (which had already built many of the transmission lines upon which Internet traffic would flow). In contrast to these high fixed costs, the marginal cost of connecting additional users to the Internet is relatively low. Use of the Internet also involves learning effects, as with any complex technology. Most significantly, the development of the Internet generates especially strong network effects. Telecommunications technologies derive their entire value from the ability to interconnect with others; a single fax machine has no utility if there are no other fax machines to receive transmissions. Similarly, the value of the Internet is largely dependent upon the number of people and information resources that are connected to it.

Not only is the Internet a technology subject to increasing returns, but it was initially designed as a technology that would be resistant to centralized control. The original engineering decisions which gave rise to this characteristic were a product of the specific economic, political, and social environment in which the Internet was created. In part, the technological characteristics of the early Internet derived from the norms of its designers and initial user community. The technology was originally the tool of a small group of engineers and academics, who were wary of bureaucracy, trusted each other, and worked well through consensus. In light of this culture, they made specific choices about the design of the technology that rendered the network resistant to efforts at centralized control (Abbate 1999).

An even more important influence on the technological configuration of the early Internet were the military imperatives for its development (Abbate 1999). The Internet has its origins in technology funded by and developed for the U.S. Department of Defense—packet switching networks designed in the early 1960s and their first large-scale implementation in the ARPANET. The rationale for packet switching technology was to design a communication network that could not be controlled from any single, centralized point, so that communications capacity could not be disabled by an enemy attack on a key portion of the network. With both the ARPANET and the later development of protocols for the Internet, survivability was the paramount goal, thus ensuring that these computer networks would not lend themselves to centralized control (Clark 1988).

The particular characteristics of the Internet that served to frustrate attempts at centralized control involve the end-to-end arguments in network design (Blumenthal and Clark 2001; Lemley and Lessig 2000; Lessig 1999). As guidelines for the design of computer networks, the end-to-end arguments state that complexity and control should be implemented at the “ends” of the network—the multiple computers and individual users that are interconnected (Saltzer et al. 1984). Meanwhile, the core of the network performs simple data transfer functions that do not require knowledge of how the ends are operating. In contrast to the telephone network, in which complex call routing is performed by a small number of centralized switching stations, the core infrastructural and computing elements of a “stupid network” like the Internet simply move packets of information indiscriminately (Isenberg 1997). Because the Internet was built around an end-to-end design, one cannot control the entire network through control of a small number of centralized nodes. Control can be exerted at the ends of the network, but as these ends multiply, controlling the entire network by controlling the ends becomes less and less feasible.

While a control-frustrating technological architecture suited the needs and preferences of the Internet’s designers and initial user community, the technology has since spread into a number of environments in which centralized control of information is a more desirable feature.

One of the most important of these major shifts involves the global diffusion of the Internet. Today, the most rapid growth of the Internet is taking place in the developing world, including a number of authoritarian regimes where standards of information control are quite different than in the United States. The leaders of these countries generally recognize the tangible benefits that the Internet has to offer, yet they worry that Internet use might pose political threats, challenge state control of economic resources, or offend local cultural sensitivities. To reap the benefits of the technology while avoiding what they see as negative ramifications, their leaders would prefer to exert greater centralized control over Internet use.

If the dynamics of Internet development mean that its control-frustrating characteristics are locked into place as it diffuses around the world, the task of authoritarian leaders is a difficult one. Without recourse to an effective architectural constraint, authorities would have no means to exert absolute control over use of the medium. Meanwhile, the economic logic of the technology's diffusion implies that there are few attractive alternatives to connecting to this control-frustrating Internet. The value of a single standardized network used by millions of people around the global far exceeds the value of any alternative network that authoritarian governments might choose to construct within their own borders.

Composite Standards, Macro-level Flexibility and the Possibilities for Internet Control

When viewed through the lens of path dependent technological development, the case for an inherently control-frustrating Internet may appear solid. This argument, however, rests upon the assumption that the architecture of the Internet is incapable of fundamental change. In this section, I delve deeper into the nature of the Internet's technological architecture, demonstrating that the composite nature of the “standard” which generates increasing returns to Internet diffusion actually gives the technology a great deal of flexibility at the macro-level. This capacity for evolution means that authoritarian leaders may be able to adapt this malleable technology for their purposes, embedding technological measures of control within the national computer networks that connect their citizens to the Internet.

To see how the architecture of the Internet might be characterized by flexibility rather than stasis, it is useful to consider the nature of the standard around which users of the Internet coordinate. In many traditional cases of path dependent technology development, coordination around a single, simple standard (e.g. the QWERTY typewriter keyboard, or the VHS format for videocassette tapes) is what generates network effects and contributes to lock-in through increasing returns. The Internet, however, involves a whole series of separate standards at different layers of the network, working together in a complex fashion to facilitate communication. The value of connecting to the Internet is not simply derived from coordination around the core TCP/IP standard as a way of exchanging data traffic. Rather, network effects in the case of the Internet are derived from coordination around the entire package—standards for e-mail, web browsing, streaming audio, encryption, and many more. At the macro-level, therefore, the Internet can be thought of as constituting a *composite* standard, with hundreds of simple standards as its constituent parts.

The composite nature of the standard involved in Internet diffusion lends great flexibility to this technology, allowing it to be adapted to meet the operating demands of new environments. At the micro-level, the individual standards for particular Internet services display a fair amount of inflexibility; once implemented and employed by millions of computers worldwide, these individuals protocols are very difficult to change.[4] At the macro-level, however, the combination of parts that make up the Internet's composite standard has changed significantly over time. The HTTP protocol for the World Wide Web, for instance, was not a part of the Internet at its origins, but it is an essential component of the Internet's composite standard today. Indeed, both e-mail and the Web—two of the Internet's most popular applications—were not originally envisioned by the Internet's creators but rather resulted from processes of informal experimentation. The Internet's macro-level flexibility has allowed it to incorporate these and other new applications as its operating environment changes over time.

Like the characteristics that rendered the Internet challenging to centralized control, the Internet's flexibility is not inherent but was rather explicitly designed into the network. Many of the same characteristics that made the Internet hard to control make it a flexible technology as well. Unlike the telephone network which was designed specifically for voice traffic, the core of the Internet was not optimized for any particular service. At the time of its creation, there was little sense of what services the Internet would need to support in the future, so the core of the network was built as a set of simple, flexible tools. Any service that conforms to the published protocols for addressing and transmitting information can be implemented at the ends of the network without altering the center. The Internet's central mechanisms simply move information indiscriminately; the core of the network does not need to know if it is transmitting packets from an e-mail, a website, streaming audio, or some as-of-yet uninvented service. Thus, the characteristics of the Internet as a whole can be altered by adding new protocols that will help the technology meet the needs of operating in new environments.

Controlling the Ends of the Internet

As the Internet spreads to authoritarian regimes around the world, its macro-level flexibility suggests that their leaders may be able to adapt this malleable technology for their own purposes. To see exactly how this might occur, it is useful to reconsider the notion of the end-to-end arguments. As principles of network design, the end-to-end arguments place users at the ends of the network. In reality, however, the Internet is much less a single network of individual users as it is a network connecting separate computer networks. Networks are interconnected through a gateway; behind the gateway, each individual network can be configured in any number of ways as long as it is compatible with the TCP/IP protocols. Conceptually, therefore, it may well make more sense to think of the Internet's component networks as its ends than to think of individual users as the outer edge of a single, seamlessly interconnected Internet.

When separate networks are conceived of as the ends of the Internet, new meaning is leant to the maxim that one can only control the Internet by controlling its ends. Exerting technological control of the Internet at the user level, in keeping with the end-to-end design principles, constitutes a quite daunting task; it would be akin to mandating that foolproof censorship software be installed on every user's computer. It is much more feasible, however, to exert control over individual networks connected to the Internet, especially where traffic passes through a single or small number of choke points.

Rather than controlling the *entire* Internet, governing authorities always attempt to control a relevant subset of Internet users. The administrators of corporate computer networks, for instance, often monitor employees' Internet usage and block certain types of non-work-related traffic. Users who have a choice of network will always be able to switch to a more liberal environment. For those with no realistic choice, however, the distinction between control of the Internet and control of a network attached to the Internet is largely irrelevant. For them, the choice is between access to a restricted Internet and access to nothing at all.

Such is the situation in many authoritarian regimes that are developing national computer networks with connections to the Internet. While in most democracies a number of individual Internet service providers (ISPs) maintain separate links to the global Internet, in authoritarian regimes all Internet users may effectively be members of a single national network. Even when there are multiple ISPs within a country, international connections to the global Internet are often channeled through a single government-controlled gateway.

Moreover, architectural constraints on the Internet at the national level can be supplemented by additional measures of technological control implemented by individual ISPs, Internet cafés, and online chat rooms. Each of these entities constitutes an additional "end" of the Internet at a level more diffuse than the national gateway but still closer to the Internet's core than the individual user. While governments may have less direct control over the technological configuration of Internet access at these levels, they can leverage their control of law and their influence over markets and norms in ways that will encourage private entities to establish their own architectural constraints on Internet use.

Technological Control of Internet Use in Authoritarian Regimes: Saudi Arabia and China

Given the political, economic, and social conditions prevailing in many authoritarian-ruled countries, one should not be surprised to find that their governments have sought to establish technological measures of control over the portions of the Internet within their borders. The governments of authoritarian regimes are typically central players in the growth of their own information infrastructures, and one would expect them to build architectures of control into their "ends" of the Internet. In the section that follows, I show how the governments of Saudi Arabia and China have sought to development national computer networks that facilitate rather than frustrate efforts at state control.

Saudi Arabia. Saudi Arabia's approach to the Internet has been strongly influenced by its conservative society, with significant public concern over pornography and material offensive to Islam, and strong support for censorship of this type of content on the Internet. In addition, Saudi Arabia is a monarchy in which the royal family is quite sensitive to criticism and dissent; it is particularly cognizant of the threat posed by overseas opposition groups like the Committee for the Defense of Legitimate Rights and the Movement for Islamic Reform in Arabia, which seek to turn public sentiment against the regime.

Because of these conditions, Saudi Arabia has moved very slowly in its approach to the Internet. The country's first connection was established in 1994, but public access was delayed until 1999 while authorities perfected their technological mechanism for Internet control. Since then, public use of the Internet has grown steadily: from 690,000 users in April 2001 to 1.46 million (or 5.7% of the population) in September 2003.[5] Saudi Arabia has chosen to permit multiple, privately-owned ISPs, but all international connections to the global Internet pass through a gateway maintained by the Internet Services Unit (ISU) of the King Abdulaziz City for Science and Technology, the Internet's governing authority in the country.

The concentrated national network structure has facilitated the technological control of Internet content, a goal about which Saudi authorities have been quite open.[6] Since the debut of public access in Saudi Arabia, all traffic to the global Internet has been filtered through a set of proxy servers managed by the ISU, aiming to block information that authorities consider socially and politically inappropriate. Market conditions have facilitated the imposition of censorship: since 1999, Saudi Arabia has outsourced the provision of censorship software to U.S.-based Secure Computing. Saudi authorities currently rely on the pre-set list of sexually-explicit sites contained in Secure Computing's SmartFilter software, which is customized with the addition of political and religious sites (Zittrain and Edelman 2002a). In addition, the ISU's website includes forms where the public can request that sites be blocked or unblocked; officials report an average of 500 block requests and 100 unblock requests per day.

China. In its approach to the Internet, China has sought a strategy which will allow it to promote widespread, market-based diffusion of the technology while still retaining governmental control. Internet growth in China has continued steadily since public access was first introduced in the mid-1990s; as of December 2004, the government estimated that there were 94 million users, or 7.2% of the population.[7] Because filtering so much traffic through a single international gateway would be nearly impossible, Internet control in China is more diffuse than in Saudi Arabia. It is difficult to ascertain the specific technological details, as China has been much less open about the configuration and extent of its censorship regime. All evidence suggests, however, that China employs multiple, overlapping layers of Internet control which have been quite effective at limiting the access of the majority of users. Zittrain and Edelman (2002b) describe a number of ways in which the architecture of the Internet in China has been modified to implement technological control. Blocking specific web pages on the basis of IP address has been the most common. In September 2002, however, authorities implemented a more sophisticated system capable of blocking pages dynamically, based on either keywords in the URL (prohibiting Google searches on specific terms, for instance) or in the actual web page requested. These methods of blocking are a step beyond previous strategies and mechanisms employed elsewhere, as they do not rely on a preexisting blacklist of prohibited websites.

At the level of the international gateway, the cornerstone of China’s Internet control has been its system of interconnecting networks. While promoting rapid proliferation of the ISPs that provide Internet access to end-users, actual connectivity to the global Internet has long been channeled through a small number of interconnecting networks with ties to government ministries or important state companies. Four interconnecting networks were initially established in 1996; the number has since grown to nine, though as the Ministry of Information Industries has licensed additional networks it has made certain that they are under effective state control (Harwit and Clark 2001). Moreover, the structure of this market is more concentrated than the number of interconnecting networks implies: the top two networks, ChinaNET and China169, jointly control 88% of international bandwidth.[8] This structure facilitates the implementation of censorship at the national level. Chase and Mulvenon (2002), for instance, report that most national-level Internet filtering is implemented by the International Connection Bureau, a set of computers belonging to ChinaNET owner China Telecom. Moreover, the major networks routinely exchange information about specific websites that they seek to block.

China has also augmented its control over Internet architecture by establishing control at the level of ISPs, Internet cafés, and chat rooms. Such points of access to the Internet number into the thousands, and most are thoroughly private entities without the same ties to the regime as the interconnecting networks. At this more diffuse level, authorities can implement an architecture of control indirectly, through their legal influence over intermediaries and the creation of a market environment in which cooperation with authorities is good business practice.

China’s Internet regulations make ISPs, Internet cafés, and chat rooms responsible for online content, and the threat of sanctions (and occasional large-scale crackdowns) have encouraged these entities to implement their own technological measures of control. It is likely that at least some of the filtering methods described by Zittrain and Edelman (2002b) are implemented by ISPs instead of (or in addition to) the interconnecting networks. For their part, many Internet cafés have chosen to install blocking software to limit what their patrons can view, and chat rooms use a technology that scans for potentially sensitive postings and sends them to a webmaster for review (Chase and Mulvenon 2002). In addition to these filtering measures, ISPs and Internet cafés have been required to implement technological architectures that facilitate government surveillance. Regulations introduced in October 2000 require ISPs to keep logs of Internet traffic for 60 days and deliver the information to authorities on request (Harwit and Clark 2001). For their part, many Internet cafés have installed software that allows public security bureaus to track user records and monitor Internet traffic remotely (Kalathil and Boas 2003).

Evidence from the cases of Saudi Arabia and China confirms the expectation that the architecture of the Internet is not inherently control-frustrating, even if this characteristic was a feature of the early Internet in the United States. Rather, the logic of end-to-end network design shows that authoritarian governments can construct national computer networks attached to the Internet in ways that facilitate technological control.

Perfect vs. Effective Control: The Importance of Institutional Constraints

While undoubtedly effective for the majority of users, the technological measures of control implemented by authoritarian regimes like Saudi Arabia and China still fall short of an absolute constraint on Internet use. Internet controls are never 100% secure; they can almost always be circumvented by determined, tech-savvy users willing to run risks and possibly pay the costs of alternative access channels. In this section, I address these inherent imperfections in technological measures of Internet control and examine the ways in which authoritarian governments have sought to supplement them by leveraging a combination of legal, normative, and market-based constraints. While perfect technological control over the Internet may never be possible, these institutional constraints are essential for establishing effective control over Internet use—a level of control that is sufficient for the political, economic, and social goals that the authoritarian leaders seek to fulfill.

Those skeptical of arguments about Internet control routinely point to the myriad ways that determined users can circumvent technological measures of control. Saudi authorities have acknowledged that many users are finding ways to access forbidden websites, often through the use of overseas proxy servers (Kalathil and Boas 2003). Wealthy Internet users who find this avenue blocked can always dial into unrestricted accounts in neighboring Bahrain—a common practice in the days before public access was permitted in Saudi Arabia. In the Chinese case, ongoing arrests of online dissidents confirm that people are successfully engaging in types of Internet use the government seeks to block. Zittrain and Edelman (2002b) and Chase and Mulvenon (2002) detail a number of ways Chinese Internet users can attempt to circumvent controls, from the use of peer-to-peer file sharing systems to entering the URLs of blocked pages in ways that may fool censorship mechanisms.

In addressing the implications of these inevitable cracks in national firewall systems, it is important to distinguish between perfect control and effective control of the Internet. Libertarian perspectives on Internet control are essentially concerned with the individual—will the government be able to prevent *me* from doing what I want to do online? For the most determined and tech-savvy users, only perfect architectural constraints will be able to control their online activity. But the perspective of authoritarian governments, or of any authority seeking to exert control over the Internet, is different. Here, the goal is almost never perfect control, attempting to thwart the evasive maneuvers of every enterprising, tech-savvy individual. Rather, authoritarian leaders seek to exert control with an external referent—control that is "good enough" with respect to any number of important objectives, including regime stability and protection of local culture. Effective control of this sort may not be able to change the behavior of the last tenth of a percent of Internet users, but this small number is rarely enough to seriously challenge the goals that most authoritarian regimes are trying to pursue.

It is in establishing and enforcing effective control over the Internet that institutional constraints on Internet use come most clearly into play. To understand the interplay of technological and institutional constraints, an economic interpretation is useful, with unrestricted Internet access thought of as a good demanded by different numbers of users depending on the price. While perfect architectural constraints, if they existed, could control the behavior of every user, institutional constraints are best seen as raising the cost of circumventing control. The cost may be literal in terms of market constraints—e.g. a satellite connection necessary to circumvent national restrictions on the Internet. In terms of law or social norms, users face the metaphorical (but still very real) costs of ostracism or punishment when they are caught.

In this economic model, most consumers are quite happy using the Internet for entertainment, online games, communication with friends, and access to officially-sanctioned news sources; they place a low value on circumventing controls, especially with regard to political information. Similarly, some percentage of users will always demand unrestricted access to the Internet even at extremely high prices; they will spend money for technology to circumvent censorship, engage in illegal political communication at the risk of punishment, and ignore disapproval from members of society who frown on lawless activity. As these costs are raised, however, demand for unrestricted Internet access shrinks. The government's goal is not to set the cost so high that demand is completely eliminated; rather, authorities seek to reduce this demand to the point of political insignificance.

Leveraging law, social norms, and the market to raise the cost of unrestricted Internet use allows for a much more effective implementation of control than architectural constraints alone. If firewalls can be circumvented with fancy technology or international phone calls, the high price of these activities helps to render this architectural constraint effective. If tech-savvy patrons of Internet cafés can configure their browsers to access pornographic or dissident websites, they will be stopped only by the ingrained knowledge that such behavior is socially unacceptable, or that café managers may be observing their Internet use and could report their transgressions to authorities.

Establishing Effective Control in Saudi Arabia and China

The cases of Saudi Arabia and China both illustrate how governments can leverage institutional constraints to establish effective control over Internet use. In Saudi Arabia, the government has found support for its censorship regime among conservative Islamist groups that are primarily concerned about pornography. Social norms against viewing material deemed offensive to Islam encourage self-censorship among users, as do legal prohibitions on accessing forbidden content and the possibility that surveillance mechanisms can identify violators. Attempts to view blocked sites are greeted with a message that all access attempts are logged;

ISPs are required to keep records on the identity of users and provide such information to authorities if requested. In addition to these legal and normative sanctions, market conditions (such as the high price of dialing into an ISP outside of the country) have also discouraged those who would seek to obtain unrestricted access to the Internet in Saudi Arabia.

In China, the use of institutional constraints on Internet access has been even more extensive, likely due to the greater challenge of exerting purely technological control over a broader and more diffuse Internet. One major way that China promotes self-censorship involves legal regulation of users. Authorities have engaged in high-profile crackdowns on various dissidents and individuals who run afoul of the regulations by engaging in politically sensitive communication. Chase and Mulvenon (2002) have offered numerous examples, from Huang Qi, who operated a website with news about the Tiananmen massacre, to members of the Falun Gong who disseminate their materials online. Sentences of several years in prison are common for such offenses, undoubtedly deterring others who might have inclinations to engage in similar activity.

Similarly, periodic crackdowns on the Internet cafés and chat rooms that allow patrons to engage in prohibited activities have encouraged these intermediaries to police their own users. In addition to implementing the technological measures of censorship and surveillance detailed above, they have added elements of human control to comply with regulations. Internet café managers tend to closely observe their users' surfing habits, especially after a series of crackdowns and closures of Internet cafés in 2001. Similarly, most chat rooms employ censors known as "big mamas" who screen postings and delete those that touch on prohibited topics. The operators of major Internet portals, who are forbidden to post information that "undermines social stability," have steered clear of anything potentially sensitive, offering primarily entertainment, sports information, and news from official sources. Even where regulations do not specifically require it, market conditions have encouraged the private sector to comply with the state's broad goals for the Internet. Doing business in China means maintaining good relations with the government. In early 2000, for instance, over 100 of China's major Internet entrepreneurs signed a pledge to promote self-discipline and encourage the "elimination of deleterious information [on] the Internet" (Kalathil and Boas 2003)

Conclusion

China and Saudi Arabia's experiences with the control of public Internet use offer a common lesson about the Internet in authoritarian regimes. Ultimately, the Internet is a tool, a medium of communication much like any other; it has no inherent political logic, no "built-in incompatibility [with] non-democratic rule" (Taubman 1998: 256). As a tool, its political impacts will depend largely on who controls the medium and in what manner they seek to use it. The Internet was initially considered an inherently control-frustrating form of communication because of features incorporated into the network by its designers. However, nothing in the technological architecture of the Internet ensured that it would remain difficult to control as it spread around the world. Rather, the architecture of the Internet is characterized by great flexibility at the macro-level, and the leaders of authoritarian regimes can take advantage of this flexibility to embed elements of control into their portions of the Internet. When leverage over Internet architecture is combined with legal, normative, and market constraints, authorities can exert effective control over the use of the Internet, preventing serious challenges to the economic and political goals that they pursue.

It is important to recognize that China and Saudi Arabia's efforts at controlling use of the Internet do not constitute mere restrictions on the diffusion of the technology within their borders. While some authoritarian regimes such as Cuba and Burma have sought to control the Internet by regulating access, China and Saudi Arabia have been enthusiastic about promoting widespread access to their national networks (though the latter did so only after perfecting its mechanism for content censorship). Rather than clamping down on Internet growth in a reactive fashion, they have sought proactive measures of control over the technology that are consistent with its rapid growth. In doing so, they are able to gain many of the economic benefits that accompany greater Internet access, as well as the improved legitimacy that may come from establishing online government services and reducing corruption and graft. While these two extreme cases of Internet control are not necessarily representative of a general trend among authoritarian regimes, they do illustrate a direction in which other countries may move in the future as they seek to emulate these successful examples of Internet control.

Indeed, there is evidence that increasing government control of the Internet is a trend not only in authoritarian regimes but among advanced industrial democracies as well. In the international security environment that has followed the terrorist attacks of September 11, 2001, the United States has placed much less emphasis on the freedom of information flow abroad and at home and has sought greater control over the Internet within its own borders. The USA PATRIOT Act legalizes a certain degree of Internet surveillance without a warrant or the establishment of probable cause. Moral concerns have also encouraged greater control: federal E-Rate funding for Internet access in public libraries depends on the implementation of filtering schemes to limit access to pornography. Finally, influential corporate interests such as the

Recording Industry Association of America (RIAA) have successfully lobbied for government crackdowns on file sharing and other technologies that could be used for copyright infringement.

In speculating about the more long-term prospects for control of the Internet, one should recall that accurately predicting the impact of a flexible technology is an inherently difficult enterprise. Given its flexibility, the specific technological characteristics of the Internet in any given environment will be largely contingent upon the political, economic, and social conditions that prevail. Moreover, the institutional constraints that influence Internet use—law, the market, and social norms—are similarly capable of change over time even when they exhibit a certain degree of stickiness. To say that China's laws and market environment or the social norms prevailing in Saudi Arabia currently support government control of Internet use does not mean that they will continue to do so fifty years hence. While it is not an automatically control-frustrating technology, a more liberal future for the Internet is certainly possible. Such an outcome, however, will depend largely on the institutional variables shaping the evolution of Internet technology and the manner in which it is used—not on any inherent characteristic of the Internet itself.

References

Abbate, Janet. 1999. *Inventing the Internet.* Cambridge, MA: MIT Press.

Arthur, W. Brian. 1994. *Increasing Returns and Path Dependence in the Economy.* Ann Arbor, MI: The University of Michigan Press.

Blumenthal, Marjory S., and David D. Clark. 2001. "Rethinking the design of the Internet: The end to end arguments vs. the brave new world." *ACM Transactions on Technology* 1.1: 70-109.

Boas, Taylor C. 2004. "Technology, Freedom, and Democracy: An Evolving Debate." *APSA-CP: Newsletter of the Organized Section in Comparative Politics of the American Political Science Association* 15.1: 18-23.

Chase, Michael S., and James C. Mulvenon. 2002. *You've Got Dissent! Chinese Dissident Use of the Internet and Beijing's Counter-Strategies.* Santa Monica, CA: RAND.

Clark, David D. 1988. "The Design Philosophy of the DARPA Internet Protocols." *Computer Communication Review* 18.4: 106-114.

David, Paul A. 1985. "Clio and the Economics of QWERTY." *The American Economic Review* 75.2: 332-337.

Franda, Marcus. 2002. *Launching Into Cyberspace: Internet Development and Politics in Five World Regions.* Boulder, CO: Lynne Rienner.

Froomkin, A. Michael. 1997. "The Internet as a Source of Regulatory Arbitrage." In *Borders in Cyberspace: Information Policy and the Global Information Infrastructure,* edited by Brian Kahin and Charles Nesson. Cambridge, MA: MIT Press.

Harwit, Eric, and Duncan Clark. 2001. "Shaping the Internet in China: Evolution of Political Control over Network Infrastructure and Content." *Asian Survey* 41.3: 377-408.

Isenberg, David. 1997. "Rise of the Stupid Network." *Computer Telephony* (August): 16-26.

Johnson, David R., and David Post. 1996. "Law and Borders: The Rise of Law in Cyberspace." *First Monday* 1.1 (May).

Kalathil, Shanthi, and Taylor C. Boas. 2003. *Open Network, Closed Regimes: The Impact of the Internet on Authoritarian Rule.* Washington, D.C.: Carnegie Endowment for International Peace.

Kedzie, Christopher R. 1997. *Communication and Democracy: Coincident Revolutions and the Emergent Dictator's Dilemma.* Santa Monica, CA: RAND.

Lemley, Mark A., and Lawrence Lessig. 2000. "The End of End-To-End: Preserving the Architecture of the Internet in the Broadband Era." The Social Science Research Network Electronic Paper Collection, <http://papers.ssrn.com/sol3/papers.cfm?abstract_id=247737>.

Lemley, Mark A., and David McGowan. 1988. "Legal implications of network economic effects." *California Law Review* 86.3: 481-611.

Lessig, Lawrence. 1999. *Code and Other Laws of Cyberspace.* New York: Basic Books.

Milner, Helen V. 2003a. "The Digital Divide: The Role of Political Institutions in Technology Diffusion." Paper presented at the annual meeting of the American Political Science Association, Philadelphia, August 28-31.

Milner, Helen V. 2003b. "The Global Spread of the Internet: The Role of International Diffusion Pressures in Technology Adoption." Paper presented at the conference "Interdependence, Diffusion, and Sovereignty," University of California, Los Angeles, March.

North, Douglass C. 1990. *Institutions, Institutional Change, and Economic Performance.* New York: Cambridge University Press.

Saltzer, J. H., D. P. Reed and D. D. Clark. 1984. “End-to-End Arguments in System Design.” *ACM Transactions in Computer Systems* 2.4: 277-288.

Taubman, Geoffry. 1998. “A Not-So World Wide Web: The Internet, China, and the Challenges to Nondemocratic Rule.” *Political Communication* 15: 255-272.

Zittrain, Jonathan, and Ben Edelman. 2002a. “Documentation of Internet Filtering in Saudi Arabia.” Berkman Center for Internet and Society, Harvard Law School. <http://cyber.law.harvard.edu/filtering/saudiarabia/>.

Zittrain, Jonathan, and Ben Edelman. 2002b. “Empirical Analysis of Internet Filtering in China.” Berkman Center for Internet and Society, Harvard Law School. <http://cyber.law.harvard.edu/filtering/china/>.

Notes

Taylor C. Boas is a Ph.D. candidate in political science at the University of California, Berkeley, and co-author of *Open Networks, Closed Regimes: The Impact of the Internet on Authoritarian Rule* (Carnegie Endowment for International Peace, 2003). His publications on this topic have also appeared in *Studies in Comparative International Development*, *APSA-CP*, *First Monday*, and *The Washington Quarterly*.

[1] Indeed, the governments of many authoritarian regimes have sought to emulate the tactics of those most successful at controlling Internet use (Kalathil and Boas 2003: 138), so China and Saudi Arabia may well serve as practical examples for others.

[2] While often conflated with the effects of high fixed and low marginal costs, network effects are a separate mechanism in that they involve increasing demand for a more widely used technology rather than a lower cost to supply that technology in the marketplace (Lemley and McGowan 1998).

[3] Arthur's fourth characteristic, adaptive expectations, is not considered here because it is not a characteristic of a technology per se but rather of its adoption process. Moreover, adaptive expectations in this context are largely a result of network effects.

[4] This does not mean, of course, that the Internet's micro-level standards are impossible to change. Indeed, there have been initiatives to alter some of the network's core protocols (Blumenthal and Clark 2001, Lemley and Lessig 2000). But this task is a more difficult one than altering the Internet by adding new functions and applications.

[5] See <http://www.isu.net.sa/surveys-&-statistics/num-users.htm>. More recent figures were not available.

[6] See the description of Saudi Arabia’s censorship regime on the ISU website, <http://www.isu.net.sa/>.

[7] See <http://www.cnnic.net.cn/>. An estimated 74 million Chinese citizens (5.7% of the population) were using the Internet as of September 2003. In relative terms, therefore, the numbers of users in China and in Saudi Arabia are quite comparable.

[8] China Internet Network Information Center, “15th Statistical Survey Report on the Internet Development in China,” January 2005, available at <http://www.cnnic.net.cn/ >.