Explorations in Cyber International Relations

Massachusetts Institute of Technology Harvard University

Control Point Analysis

David D. Clark

Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology

September 10, 2012

This material is based on work supported by the U.S. Office of Naval Research, Grant No. N00014-09-1-0597. Any opinions, findings, conclusions or recommendations therein are those of the author(s) and do not necessarily reflect the views of the Office of Naval Research.

Control Point Analysis

David Clark

MIT CSAIL

TPRC submission version 2.2 of September 10, 2012

This work is funded by the Office of Naval Research under award number N00014-09-1-0597 as part of the DoD Minerva program. Any opinions, findings, and conclusions or recommendations expressed in this email are those of the author and do not necessarily reflect the views of the Office of Naval Research.

Introduction

As the Internet becomes more and more embedded in every sector of society, more and more actors have become concerned with its character, now and in the future. The private sector actors, such as Internet Service Providers or ISPs, are motivated by profits as they shape and evolve the Internet. The public sector is driven by a range of objectives: access and uptake, competition policy, regime stability, policies with regard to controlling access to classes of content, and the like. The range of actions open to governments to shape the Internet are traditional and well-understood, including law and regulation, procurement, investment in research and development, participation in the standards process and more diffuse forms of leadership. But these actions do not *directly* shape the Internet. They bear on the actors that in turn have direct influence over the Internet and what happens there. Thus, as part of any conversation about the shaping of the Internet, there is a narrower question that must be answered: given the Internet as it is today, who are the actors that can exercise direct control over how it works, what options for control do they actually have, and how can they in turn be influenced?

These questions require an understanding of the Internet as a technology, which can be a bit of a daunting task. A technical description of a system like the Internet usually begins with its modularity (e.g. layers and regions), and the functions and formats of its protocols. These sorts of descriptions are often not of much use when describing a system to a non-technical listener—the mass of unfamiliar details masks any insights about the *implications* of the design with respect to issues such as economics or the relative power of various actors to influence the operation of the system.

This paper describes an informal method I call *control point analysis* that can be used to capture and understand the power relationships created by specific design decisions surrounding the Internet. In particular, control point analysis focuses on the question of finding the locus of power and control implied by the design—is power centralized (and if so, to what actor) or diffuse? Does the design create points of control or avoid them? A useful conversation across disciplines must begin with a method of extracting and cataloging the important implications of the design without first getting lost in the technical details of the design. Control point analysis is a possible method for doing this.

Does technical design matter? Who controls cyberspace?

Different designs for cyberspace can have major implications for the balance of power among the various interested actors. This consequence may not be obvious to all network designers, but it has long been very clear, at least to some. In the 1970's, there was a substantial debate between advocates of two sorts of network, called "datagram" and "virtual circuit". Datagram networks have a simpler core, with more functions shifted to the hosts at the edge. Virtual circuit network have more function in the core of the net, and thus more power and control shifted to the network operator. The Internet is a datagram network; the ARPAnet was more a virtual circuit network, and the data network standard developed by the telephone industry, Asynchronous Transfer Mode, or ATM, is a virtual circuit network.

One of the vocal advocates of the datagram approach was Louis Pouzin, who was building a datagram network called Cyclades in France at the same time that the Internet was being first built. In 1976, he published a paper with the following conclusion[1]:

> *The controversy DG vs. VC in public packet networks should be placed in its proper context.*
>
> *First, it is a technical issue, where each side has arguments. It is hard to tell objectively what a balanced opinion should be, since there is no unbiased expert. This paper argues in favor of DG's, but the author does not pretend being unbiased. Even if no compromise could be found, the implications would be limited to some additional cost in hardware and software at the network interface. So much resources are already wasted in computing and communications that the end result may not be affected dramatically.*
>
> *Second, the political significance of the controversy is much more fundamental, as it signals ambushes in a power struggle between carriers and computer industry. Everyone knows that in the end, it means IBM vs. Telecommunications, through mercenaries. It may be tempting for some governments to let their carrier monopolize the data processing market, as a way to control IBM. What may happen, is that they fail in checking IBM but succeed in destroying smaller industries. Another possible outcome is underdevelopment, as for the telephone. It looks as if we may need some sort of peacemaker to draw up boundary lines before we call get it trouble.*

In contrast to the Internet, Pouzin's Cyclades network was not ultimately successful. Its failure is often (if speculatively) attributed to the hostility and resistance of the French PTT.

---

[1] Pouzin, L. 1976. Virtual circuits vs. datagrams: technical and political problems. In *Proceedings of the June 7-10, 1976, National Computer Conference and Exposition* (New York, New York, June 07 - 10, 1976). AFIPS '76. ACM, New York, NY, 483-494. DOI= http://doi.acm.org/10.1145/1499799.1499870

Control point analysis of the current Internet

Control point analysis, as I use the term here, is not a highly formalized methodology. It is an approach to help think in a methodical way about the design of a system from a particular perspective—that of determining which actors obtain power, economic or otherwise, by virtue of control over key components of the system.

The end-point of this analysis is a catalog of all the relevant actors in the ecosystem, and the forms of control and power that they exercise. But some methodical process is appropriate to generate this list.

Control point analysis proceeds by listing the steps of common actions (for example, retrieving a Web page), and asking at each step if one has encountered a significant point of control[2]. One must be methodical and complete as one catalogs the steps, but in this way one can identify control points in the ecosystem that might have been overlooked in the initial catalog of the technical system and its parts.

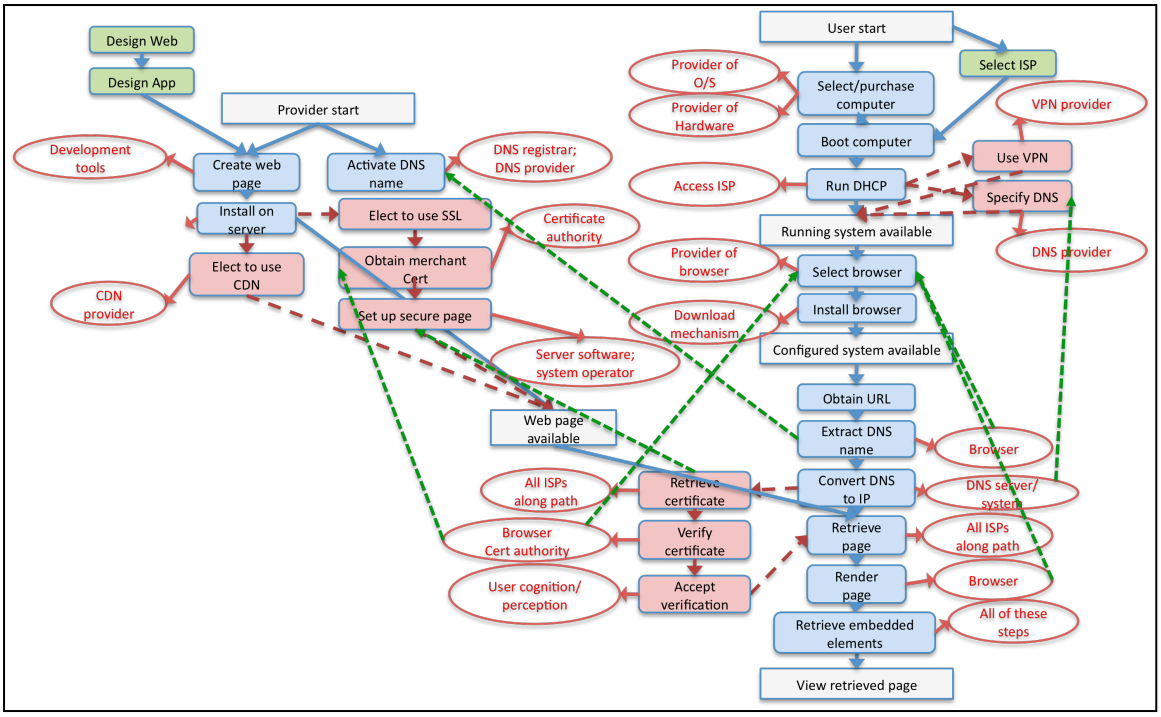

Figure 1: Steps in the retrieval and viewing of a web page. Blue arrows indicate the normal sequence of steps. Green arrows capture dependencies on prior steps. Red ovals catalog the actor(s) that have control of the outcome of each step. Green boxes are parts of the diagram that are not elaborated in this paper.

[2] A similar sort of diagram can be found in Koponen, T., S. Shenker, et al. (2011). "Architecting for innovation." SIGCOMM Comput. Commun. Rev. **41**(3): 24-36. where it is called information flow mapping.

Retrieving a Web page

Figure 1 illustrates the steps that lead up to the retrieval of a Web page. At this scale, it is perhaps most useful as an eye-chart, and is presented in this form primarily to illustrate the scope and complexity of the process. For easier viewing, figures 2, 3 and 4 break this figure down into three parts: preparation of the computer, preparation of the web page, and the actual retrieval of the web page.

In figure 2, the reader will note that the sequence of steps has been taken back to "the very beginning": step 1 is "select and purchase computer". This may seem extreme, but in fact it helps to remind us that it is not just the technical features of the system that matter in determining whether we can successfully accomplish our goal or whether there are actors with the power to disrupt or manipulate what we are doing toward an undesirable outcome.

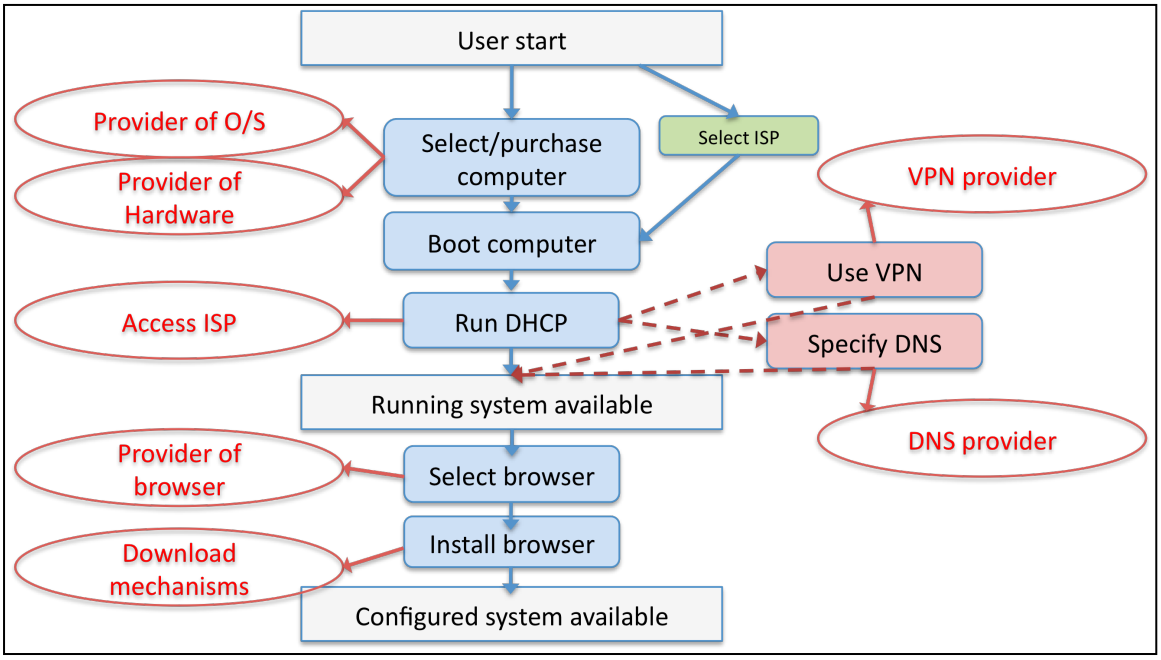

Figure 2: The initial steps that prepare the computer for use on the Internet. Optional steps (often taken to increase security) are in dark pink.

Each step in the process has been annotated (in a red oval) with the actors that have influence over the action—actors that can cause it to be successful or disrupted. The picture is thus not just about identifying actors, but it is a map of security vulnerabilities—one sort of power that an actor can hold is the power to disrupt or attack. So one way to read the annotations is that at each step, success must either depend on the trustworthy nature of the actor or the constraints within which the actor sits that limit his ability or reduce his motivation to disrupt.

- Thus, the first step where the user purchases a computer reminds us that we depend both on the hardware and software to work as expected. Concerns about corruption of the hardware and software supply chain, which might lead to malicious hardware or pre-installed malware, are captured here.

- After a computer is booted and started up, it normally runs a protocol called Dynamic Host Configuration Protocol (DHCP). In this step, which happens automatically if the computer is connected to the Internet; the access ISP gives the computer its IP address, the address of a machine (a router) that is the computer's path into the Internet, and the address of a DNS server. The Domain Name System (the DNS) is the system that translates names (for example, names like www.example.com) into actual IP addresses. The DNS server is the starting point for this conversion service. Since the DHCP protocol specifies the DNS server to be used, the user must trust the access ISP to provide a trustworthy DNS server, unless the user takes the optional steps of manually configuring a DNS server or opening a VPN so as to effectively connect to a different access ISP.
- The user may optionally choose to download and install the browser of his choice, as opposed to using the browser that came with the system. Today, there are open source browsers (Firefox), and more closed browser like Internet Explorer.

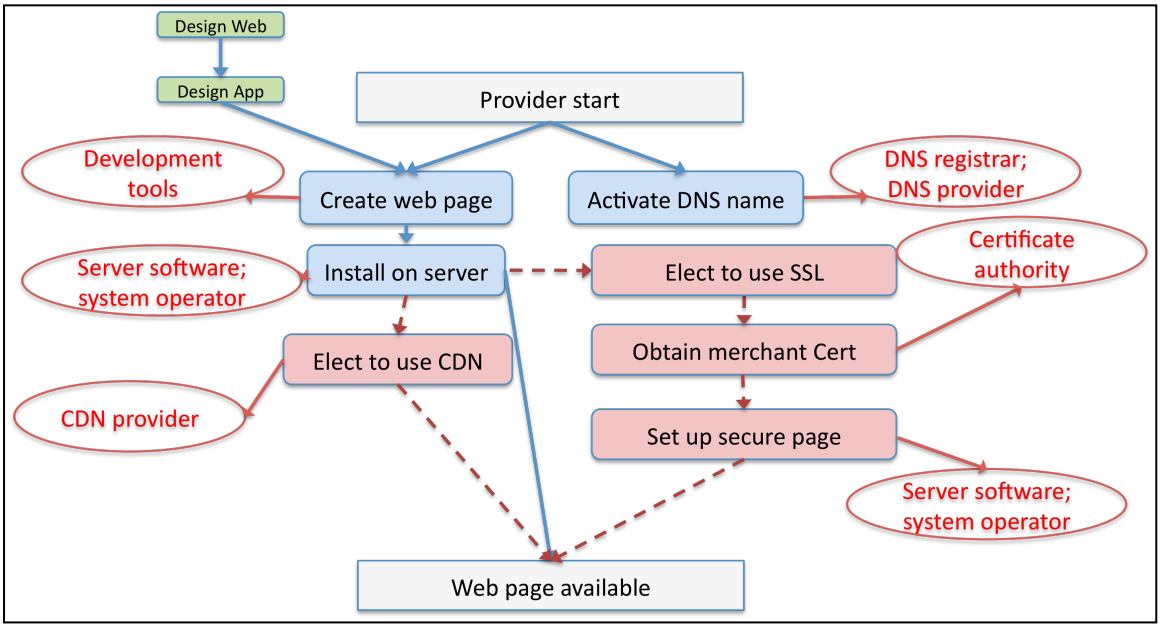

Figure 3: Steps taken by a content provider to ready a page for viewing on the Internet.

Figure 3 captures the steps that a content provider takes to make a web page available on the Web.

- The first (and obvious) step is to make the page. The correct outcome depends on the development tools, as well (of course) as on the skill and attention of the developer.
- The next step is to specify the URL for the page, which requires that the Domain Name of the server be known and included in the URL. For example, if the URL were http://www.example.com/really/cool/page, then the DNS name is

www.example.com. This requires that the DNS name (example.com) be obtained by purchasing it from one of the many providers of names, and also requires that the relevant DNS servers be configured to "resolve" the name—that is, return the address of the server when the name is looked up in the DNS.

- The final required step is to install the page on the selected server. Both the software that implements the server and the human operators that maintain the server have the ability to influence this step. In particular, poor attention by human operators is the source of many security vulnerabilities and thus loss of control to attackers.
- In addition to these basic steps, the provider may choose to improve the availability of the page by contracting with a Content Delivery Network (a CDN) to replicate the page at points across the Internet, which creates a dependency on the CDN provider.
- The provider may also choose to use secure Web protocols (Secure Socket Layer, or SSL) to enhance the security of the anticipated downloads. To do this, the provider must obtain a *merchant certificate* from a *certificate authority* (a CA)—essentially a signed verification that the provider of the page is who it claims to be. Encryption tools are used to create and validate these certificates.

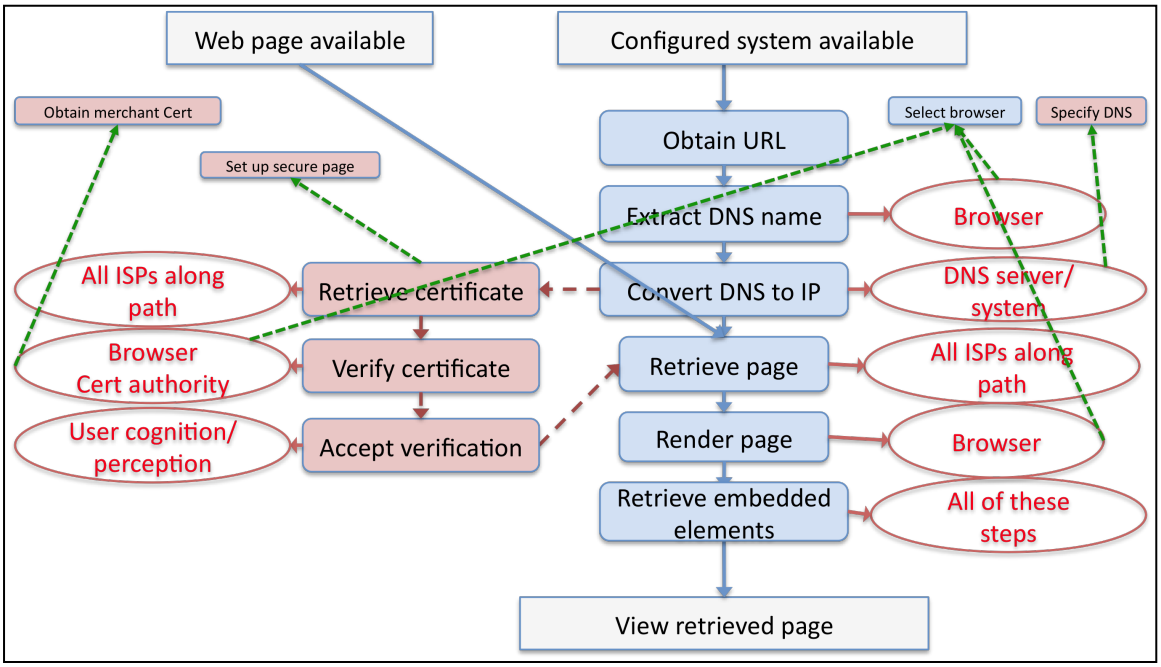

Figure 4: The steps that make up the actual retrieval and viewing of the page. Green arrows to smaller boxes (which are from previous figures) capture various dependence on the prior steps.

Once the user has a running system, and once the provider of the web page creates the page and makes it available on the Web, the actual steps of retrieving the page can occur, as illustrated in figure 4:

- The user acquires a URL by some means, perhaps from a search, by typing it in or by clicking on a link in another web page. This action is carried out using the browser software that either came with the system or was selected as one of the steps in figure 2. The provider of the browser is a relevant actor with power in the ecosystem. There have been claims in the past that certain browsers would not allow the user to use certain URLs, but this does not seem to be a major concern today. The browser presents other options for control; see below.
- The DNS name in the URL must be extracted and translated into the IP address of a server, which requires the use of the Domain Name System, or DNS. As described above, the system begins the process of translation by connecting to the DNS server that was specified in the steps in figure 2, either provided by the access ISP or chosen by an optional action taken by the user. The DNS protocols and interfaces are specified in open IETF standards, and for a long time were not seen as an important point of control. However, the DNS system itself is highly decentralized, with most ISPs operating a server for their clients. For this reason, each ISP (or other service provider, such as a hotel or hot-spot) has a very high degree of control over what address is returned in response to a DNS lookup. Many mis-directions occur in practice today using this point of control, and DNS servers have also been the target of attackers who install mis-directions of their own. Secure DNS (DNSSEC) provides technical tools to prevent a benign server from being misled by a malicious server, and at best can allow a user to tell that he has received invalid information.
- Assuming that the DNS has returned an IP address, the browser opens a connection (in tech-speak, a TCP connection) to that address. The routers at each hop along the path to the server look up the preferred path toward that address and forward the packet. They thus have absolute control over whether the client can reach the server. If the router has no route computed to that server, no connection can occur. (This outcome does not normally arise, but can during transient outages.) More significantly, if the router deliberately mis-directs the packet, or changes the destination address in the packet, the packet will arrive at the wrong server[3]. Secure BGP, now being pushed for deployment, provides tools to prevent a benign region of the net from being misled by a malicious region.
- If the web site uses secure protocols (signaled by the prefix HTTPS rather than HTTP at the beginning of the URL), the server will return to the browser a *certificate* attesting to the identity of the server. The certificate is signed or

[3] This sort of mis-direction may seem unlikely, but it (or a related mis-direction involving the DNS) occurs all the time. It is a common experience to open a browser window in a hotel or "hot-spot", and attempt a connection to some page, only to get a page instead inviting the user to pay a fee for access. This happens only because of some intentional, if mostly benign, mis-direction occurring within the hotel/hot-spot system.

validated by one of a number of *certificate authorities* or CAs, as was set up in the steps of figure 3. There are a number of important points of control surrounding these CAs. Different CAs may exercise different degrees of care before signing a certificate. But more interesting, all browsers today have built into them a list of CAs that are deemed trustworthy, which gives the browser designers a significant measure of control[4].

The use of certificates can detect some of the forms of mis-direction that occur at the DNS or router level. That is, if the DNS or router have mis-directed the browser, this can be detected. The consequence is that the browser will raise some sort of alert or alarm to the user. However, most users have no idea what to make of these alarms, and often proceed ("click through") to connect to the wrong server, to their peril.

- The web site then finds the stored content associated with the URL (or prepares the content dynamically if this is what is required) and returns this content to the browser over the network. The content may have "embedded content": URLs of further pages that are to be retrieved by the browser and displayed as part of the page. This process requires the browser to repeat all of the above steps for each of those embedded content links. Each of them may be subjected to mis-direction by the same points of control. Some of the embedded content may be hosted on other servers. The most common example of this is advertizing. Advertizing raises specific risks aside from annoyance—since malware can be hidden in innocent-looking web pages, any web server that includes third-party ads on its web site must trust that the site generating the ads is trustworthy and has not been infiltrated. There is no way for the original web server to check.
- It might seem that the site hosting the content has (as one would expect) ultimate control over the format and content of the page. However, for a number of reasons, this is not correct. Since the ISPs along the path from the server to the browser control the routing, any of them can redirect the returning web page to an intermediate node that performs arbitrary modifications to it before sending it on to the browser. Unless secure connections are used (SSL, as described above), the power of the ISP to control routing gives it the ultimate control. Examples of modifications by ISPs that have occurred today include reformatting the page to fit onto the small display of a mobile device (which seems somewhat benign) and finding embedded content URLs and replacing them—for example replacing the ads selected by the web server with ads selected by the ISP. This behavior is not normally considered benign.

This sort of control point analysis reveals that the Internet, although sometimes described by its creators as "simple", contains a rich mix of points of control, and a range of design principles that different actors use to "blunt the instruments of control" by other actors.

---

[4] If a user encounters a certificate signed by a CA that is not included in the list included in the browser, as for example with certificates issued by MIT, strange error messages arise, and the user is instructed to take inexplicable actions, such as "downloading a new authority certificate". This outcome degrades usability.

Encrypting content is an obvious example of an application-level design option to protect the content from control by others—an ISP cannot change what is signed, and cannot see what is encrypted. Other approaches used to blunt the controls of the ISPs include the use of Virtual Private Networks (VPNs), a common tool for travelers using ISPs, hotels and hot-spots they do not trust. A VPN provides an encrypted path over which it then sends all traffic from the client host back to a trustworthy relay point (e.g. at the traveler's corporate headquarters), where the traffic is then injected into the Internet as if it originated there. The IP addresses that the client machine is trying to reach are encrypted until they reach the trustworthy relay site, so intermediate untrustworthy routers cannot see them.

Table 1 provides a summary of these various points: it is a list of the actors from the three figures with two further notations: the sorts of failures and disruptions that the actor in question can inject, and the range of constraints that limit the actions of this actor.

Table 1: List of steps that comprise the retrieval of a Web page

Step	**Optional step**	**Controlling actor(s)**	**Examples of problems**	**Constraints on abuse**
USER steps				
Purchase computer		Hardware designer/manf	Corrupted supply chain, DRM	Lost reputation and market share
		Software (OS) provider	Buggy code	Lost reputation and market share
Select ISP				
Boot computer				
Run DHCP		Access ISP	NAT, address of untrustworthy DNS,	Variable—regulation, reputation, no constraints
	Initiate VPN	Access ISP, VPN provider	Access ISP can block VPN	VPN: Persistent relationship with provider
	Select alternative DNS server	Access ISP, DNS provider	Access ISP can block access to remote DNS	DNS provider: reputation.
Running machine				
	Select/download preferred browser	OS, download steps (as described here), maker of original and preferred browser	Some browsers have refused to download other browsers (in past).	Loss of reputation
Configured machine				

WEB PROVIDER steps				
Design Web		Standards bodies		
Design application		Standards bodies, innovators	Consumer tracking	
Create web page		Development tools		Loss of reputation
Activate DNS name		Domain registrars, sellers	Name disputes, release of PII	Loss of reputation, ICANN intervention (?)
Install web page on server	Server software (e.g. Apache), SysOps, Provider hosting server.	Poor server configuration and lack of patches lead to penetration and installation of malware for subsequent download.		
	Utilization of CDN	CDN provider		Persistent business relationship
	Elect to use SSL			
	Obtain merchant cert	Certificate authority (CA)	Lax attention leads to penetration of CA and creation of false certs.	Loss of reputation, business. Lawsuits. (?)
USER steps				
Obtain URL		Depends on source	Phishing attacks	Cognition and perception of user
Extract DNS name		Browser		
Get IP address of server		DNS server selected/provided above	Mis-direction to wrong IP address	Highly variable constraints depending on context
	If SSL, retrieve cert	All ISPs along path		
	Verify cert	Browser, CA	Corruption of CA	
	Accept result of verification	User downloading page	"Look-alike" names	Cognition and perception of user
Attempt to download page from server		All ISPs along path	Adverse outcomes. Without SLL: delivery of wrong/malicious version of page.	

			With SSL: unambiguous failure if mechanisms run properly	
Render page		Browser		
View page				

Given the range of actors that can exercise control over the attempt to download a web page, a variety of outcomes can occur in practice (and do occur) when a download is attempted. Table 2 summarizes the range of outcomes that a user can encounter in practice today.

Table 2: summary of outcomes when a web page is retrieved

Primary outcomes		Examples
Intended	Correct page viewed	
Unintended	Wrong cert -> failure	
	Benign delivery of wrong page	Hot-spot login page
	Malicious delivery of wrong page	Phishing, ISP redirection, notification of blocked access.
	Attacker in the middle	Modification of page; capture of user info
	No apparent response	ISP blocking, server down; network failure
Secondary outcomes		
	Intended side-effects	Delivery of cookies; tracking by provider; release of behavioral information to third parties
	Unintended side-effects	Delivery of malware; corruption of browser/system; theft of user information

Why does the Internet work? Some observations

Reputation

Looking down the right-hand column of Table 1, one point that stands out is that the constraint that disciplines many of the actors is a fear of loss of reputation. For private sector actors, loss of reputation translates into loss of business and revenue. For public sector providers, the loss of reputation can have other consequences. For some actors, the loss of reputation as a result of disreputable Internet behavior is minimal. Very few people walk out of a hotel because of the manipulation of its Internet service, although there are some who do not return. Some of the constraints may be codified as law, so that legal retribution might result from misbehavior, but in many cases, the discipline is less

formal. In fact, what this table reflects is that norms and expectations shape proper behavior by actors such as ISPs.

Technical mechanisms

The role of technology is very specific. Secure connections (SSL), secure BGP or secure DNS do not ensure correct operation. They can protect benign and trustworthy regions of the Internet from being controlled and disrupted by untrustworthy regions, but the best they can do in general (if they work as desired) is to detect incorrect operation and give the user an unambiguous signal of failure. It is up to the user to avoid or bypass the mis-functioning (or malicious) component, or somehow demand proper behavior. Tools such as manual configuration of DNS and VPNs serve as bypass tools to reduce the influence of an untrustworthy access ISP.

This analysis suggests a more general principle for network design, if the goal is operation in the presence of untrustworthy elements. What network mechanisms must do is translate arbitrary intervention by various actors into a clear signal to the user of a problem. What the network architecture should do is give the user the ability to "route around" as many actors as possible. For example, multi-homing and user selection of routes at a suitable level are means to bypass ISPs that are exercising undesired control. Of course, some secondary actors (e.g. nation states, rights holders and the like) want to impose their controls on users, and their goal is to prevent bypass or "routing around". This is an essential tussle of control.

The role of choice as a discipline

If the user is to "route around" a misbehaving actor, the design of the system must give the user that degree of choice. The tussle of control is often thus a tussle over who controls the choice. Examples above include which ISP to use, which DNS to use, which browser to use, and there are more subtle and complex choices that are embedded in the control picture. Different designs give control of these choices to different actors, so one must review the catalog of actors to see, in the context of a specific design option, which actors have control over which choices.

Cataloging the major actors

If we look at the major actors that make up the Internet, we see that different actors hold different points of control with different powers—if there is engagement or tussle among these actors using their powers, it is asymmetric engagement.

- Internet Service Providers (ISPs) build and operate regions of the network. Within their regions, they control topology and completion of connections. (e.g. who talks to whom under what circumstances.) There is no monolithic Internet, but different parts controlled by different ISPs that may be trusted (or not) to different extents. Examples of control include physical topology and routing (making sure traffic actually passes through a firewall or point of censorship). They exercise ultimate control: if they do not forward packets, the operation fails.
- Application designers, by their decisions, define the available patterns of communication, and thus shape the consequences of intervention by various other

actors. One could analyze a range of applications, just as we did the web. One could diagram sending an email, a VoIP call, or a music sharing protocol. In each case, what we would see is that the design of the application shapes the overall patterns of success and failure. For example, different placement of services and servers (and different degrees of decentralization) change the options for bypassing undesirable actors. Different uses of encryption determine what is revealed or concealed in the messages being communicated. Encryption can occur at different levels in the system (link, VPN, end-to-end or application), but only the application level can discriminate among different parts of the communicated information, encrypting some but revealing other parts. For example, the design of email protocols reveals the headers of the email, even if the body is encrypted. This design decision actually reveals a great deal of information, but at the same time permits staged delivery, which in turn allows for the "out-sourcing" of virus checking and spam filtering.

- Users and their end-node computers control the initiation of activity. To the extent they have choice in the selection of service providers, they can use that choice as a discipline to select for trustworthy providers.
- The operating system designer provides the platform for execution of end-node software. While some platforms today are more open (Linux) and some are more controlled by their owners (e.g. Windows), most operating systems today are viewed as raising few fears of explicit exercise of control. However, some do exercise considerable control: the iPad operating system and browser, for example, will not render Flash elements in web pages.
- The DNS system and the distributed servers that implement it play a critical role in translating names (e.g. URLs) into IP addresses. The system is highly decentralized, with different regions operating under different constraints, and being more or less worthy of trust. Secure DNS (DNSSEC) can (if fully implemented and deployed) help ensure that the results of a DNS query are valid.

Other actors

Control point analysis reminds us of the actors that can intervene directly in the operation of the network. There are tiers of actors behind them that can influence the behavior of this first tier of actors in many ways, from offering economic incentives to writing standards and passing laws and regulation (or encouraging the passing of such laws).

Private sector actors

Some of the actors in this category are established industries that have been strongly affected by the Internet. In most cases, they do not exercise significant direct power of points of control in cyberspace, but many of them have demonstrated considerable ability to shape cyberspace indirectly, by shaping legislation, regulation, standards and public policy.

- Telephone companies and their suppliers
- The music industry
- Radio
- The video/movie/TV industry

- "Brick and mortar" merchants of various sorts, such as book-sellers
- The "print media" industries: newspapers, magazines, etc.
- Publishing generally
- The advertizing industry
- Gambling, pornography and other marginal social activities.

Other actors (or activities) in this category include those that have emerged as a result of Internet/Cyberspace

- Computer games, massive multi-player games, virtual worlds.
- Online auctions (eBay, etc.)
- Search providers (Google, etc.)

Governments

Governments, as the traditional actors on the stage of domestic policy and international relations, are clearly important in this analysis. Again, they do not normally exercise direct control over cyberspace, but can exercise great influence by their ability to influence other actors using regulation, legislation, investment (procurement and research) and standards.

International governance organizations

This special class of NGOs includes standards bodies, such as:

- IETF
- ITU

These actors clearly exercise great power, and thus are the targets in turn of other actors that want to exercise indirect influence.

The category also includes actors concerned with governance of cyberspace, such as:

- Internet Governance Forum
- ICANN

Illegitimates

This category includes classic crime categories such as confidence games, extortion, fraud, identity theft, etc. It also includes emerging state and non-state actors using tools such as terrorism.

NGOs

Individuals

One hypothesis about the current world is that the Internet (and cyberspace taken broadly) seems to have shifted the balance of power toward certain actors, such as NGOs and the individual.

Norms

Many of the behaviors that we expect of the various actors—ISP, DNS providers and so on—are not defined by precise standards or laws. They are commonly held expectations—norms of behavior. Part of the struggle today with the future of the Internet

is to come to an understanding of what these norms are. Attempts to codify what is generally understood are often failures.

The DNS provides a good example of a norms-based domain where attempting to codify the norms is very difficult. As a starting point, one might propose the following as the norm that should define the operation of the DNS:

> *The owner of a DNS name should have sole control over what address is returned by the DNS system when that name is looked up.*

In other words, Google, and no other actor, should be able to specify what address is returned when one queries a DNS name such as www.google.com.

This seems like a nice norm, and if everyone obeyed it, it would seem to eliminate many of the mis-directions and abuses that occur today. But if one were to try to codify this norm, it would almost certainly trigger great pushback. One reason has to do with what happens at hotels, hot-spots and the like. When one first connects and attempts to go to a web page, the page that is actually returned is the login page (payment page) of the provider. This redirection is implemented by either modifying the result of the DNS lookup, or modifying the routing. This sort of "messing with" the DNS would violate that norm. And objections would not only come from those sorts of providers, it would come from governments. Governments, in their push to control access to illegal content (defined by the various laws of the various lands) have been looking for mechanisms of control, and the DNS system is an obvious target. Many governments, including the U.S[5], have considered mandating that the DNS service providers return the "wrong" answer for domains that have been found to host unacceptable content.

One could consider a modified form of the norm, as follows:

> *The owner of a DNS name should have sole control over what address is returned by the DNS system when that name is looked up, except to the extent that the law of the land specifies otherwise.*

This version of the norm would allow for at least the intervention of the state, as described above, but would probably make people very uncomfortable if it were actually written down and debated, because it would seem to legitimize the actions of more repressive states, which filter vast amounts of content, doing so, of course, consistent with the law of their land. Norms are tricky things.

One could have similar discussions about norms concerning the level of care to be expected of Web hosting services, operators of the Internet routing system, the level of training and care to be expected of normal Internet users, under what circumstances web pages can be modified as they transferred from the server to the user, or blocking of VPNs. To varying degrees, all these norms prove slippery if one tries to nail them down.

Control and observation

The term we chose for this process, control point analysis, suggests that the only objective of the process is to explore options for active control—options for manipulation

[5] As illustrated by the recently proposed but contested Stop Online Piracy Act (SOPA).

or modification of the intended task. But there is another dimension to the analysis, which is to ask, at each point of control, what options the actor has to observe what is being done—what we might call "observation point analysis", if we wanted a more complex phrase to describe this method.

Control is perhaps easier to think about, because control is an active intervention. Presumably, it has visible and immediate consequences. But passive observation—spying, monitoring, or whatever—does not manifest in immediate consequences. Its consequences are more diffuse, and can occur later. They could include behavioral profiling, revelation of personal information, punishment for unacceptable usage, and the like—a broad range of outcomes. However, for each point of control, one can construct a table similar to our control table, but focused on what is revealed, and what limits, if any, govern that revelation and its consequences. And many of the methods that the user can use to thwart unwelcome control also serve to thwart unwelcome observation. VPNs, because they encrypt what is being sent, limit what can be observed. In fact, it is this obscuration that helps thwart control—since the observer cannot see all the various things the user may be doing, his tools of control are blunted: all he can do is interfere indiscriminately, which would be an ineffective version of control in most (but not all) cases.

Controlling the Internet—four case studies

Given the multitude of options for control, it might seem that an organized actor could find some way to exercise control if that were his intention. Case studies may help illustrate the space of contention. In figures 5-9 we illustrate the points of influence exercised by five interesting actors: a typical U.S. ISP, the U.S. government, owners of content who are concerned with piracy, Google, and the Chinese government.

A typical U.S. ISP

The first case, the ISP, is illustrated in Figure 5 and was discussed earlier in the paper; as an actor with direct access to control points in the flow of packets, its influence is significant, but is concentrated in the steps where packets are being sent, or (because of the ISP's control of the DNS servers) when a name is resolved into an IP address.

The U.S. government

In Figure 6, I add the U.S. government as an actor. It does not act directly on the Internet, but acts indirectly, by means of law and regulation bearing on other primary actors, such as ISPs and content hosting sites. It also exercises a more long-term and indirect influence through its procurement, funding of research and the like. These are not illustrated in the figure.

Content owners

The music and movie industries have been trying to control the flow of unauthorized (pirated) material over the Internet. Since they do not have direct access to any of these points of control, they have been forced to work indirectly, often using laws they negotiated for the purpose. Figure 7 illustrates their points of control.

- Select/purchase computer: Content owners have worked with the computer industry to add mechanisms to computers to regulate how protected content can be used. In general, these mechanisms are called Digital Rights Management, or DRM.
- Select ISP: In the U.S., content owners can demand of ISPs that they reveal the identity of a user at a particular IP address. In some countries, they have persuaded ISPs to ban users that traffic in pirated material.
- Convert DNS to IP: The content owners worked with supporters in the government to propose a law (the Stop Online Piracy Act or SOPA), which would have authorized government to order that DNS providers return the "wrong" address of sites hosing infringing material. (At the present time, this proposal is not being pushed forward.)
- Install on server: Content owners can demand of hosting sites that they take down unauthorized copyright materials, under the terms of the Digital Millennium Copyright Act.
- Retrieve page: They can become users of the system, observe which sites are hosting the content, and demand of the access ISPs that they disclose who the owner of the site is. (This option illustrates the "observation point" aspect of this paper.)
- Provider start: They can bring lawsuits against providers of infringing content, once the ISPs have provided their identity.

In general, they have only two options—enlist the aid of an actor that has direct access to one of the points of control, or work to have the design of the system changed so that there are new options for control—in essence redraw the control point picture. The latter is hard, because they do not control the design. This paper just looked at one application—downloading a web page. But every application has its own control point analysis, and the users interested in sharing unauthorized content try to deflect intervention by designing new applications (specifically peer-to-peer systems) that try to avoid points of control.

Google

In Figure 8, I single out Google as an example of an important and powerful private-sector player, specifically because of the many actions they have taken to shape the Internet experience.

- Select/purchase computer: Google has developed and made available to smart phone manufactures the Android operating system, in order to increase choice in the consumer marketplace.
- Select Browser: Google has developed a browser called Chrome, which is available for free download. Chrome attempts to offer enhanced features for web browsing, and enhanced security for Google downloads.

- Obtain URL: Google, of course, is the major search engine in many parts of the world. As a point of control, they do customize and in some cases filter search results in various parts of the world.
- Specify DNS: Google makes available a DNS server that anyone can connect to, in order to avoid servers that may be returning inappropriate answers.
- Create web page: Not only does Google return search results, it is a provider of Youtube, one of the most popular sites on the Web.
- Elect to use CDN: Google has built its own Content Delivery Network, with global reach and direct connection to many consumer-facing ISPs.
- Retrieve page: Because Google has its own network with global reach, in many cases content downloaded from Google crosses only two ISPs, the access ISP of the consumer and the Google network. This configuration reduces the number of ISPs that might otherwise be in a position to manipulate the transfer.

China

Typically, governments do not have direct access to most of the control points in the diagram, and must act indirectly. However, in the case of China its leverage over some of those actors is considerable; for example the larger ISPs in China are essentially state-owned. Given the power of the Chinese government, their role seems almost that of a direct controller.

- User start: China has arrested users who are sources of unacceptable content on the net.
- Select O/S: China developed filtering software (Green Dam) that was to be installed in all Windows systems. This effort only partially succeeded; the software is on computers in Internet cafes, but not necessarily on personal computers.
- Select ISP: China requires that all ISPs, including mobile hot-spots, obtain and retain the identity of each user.
- Use VPN: China regularly blocks protocols such as VPNs and more sophisticated bypass software such as TOR, either by blocking the protocol or the destination port number.
- Run DHCP: China can impose restrictions on which DNS server is used.
- Obtain URL: China imposes requirement on providers of search tools to remove unacceptable content from the search results.
- Convert DNS to IP: By their control of the DNS system, they can return incorrect IP addresses for blocked content, or return no answer at all.
- Retrieve page: China instructs its ISPs to control routes, especially at their borders, block access to certain applications (e.g. Facebook, Google, Twitter, and so on), block access to specific websites, block circumvention protocols, and use deep packet inspection (DPI) to look for specific keywords in the packets and

terminate the connection. Attempts to reach an IP address that returned a sensitive keyword may be blocked for some period.

- Design applications: China has blocked many popular web applications such as Facebook, Twitter, Google, eBay, PayPal, Skype or Youtube. They have been replaced by Chinese alternatives, which are designed consistent with Chinese language and culture, but which also allow nuanced control of content.
- Provider start: Providers of certain sorts of web content, such as online forums or audio and video services must register with the government.
- Install on server: China has constructed a complex socio-technical framework to detect unacceptable content and mandate its removal or modification.

The final outcome?

There is no final outcome. All of these actors contend to shape the Internet experience as they would prefer. It would seem that the tussle over control is ongoing, with no lasting victory for any side. One actor designs an application, other actors hunt for points of control, others design mitigations to the controls and so on. This assessment should not be a surprise.

Each oval that illustrates a control point can have a very complex story behind it—a story that might itself take a diagram or a paper to capture. For example, the step "retrieve page", which depends on "every ISP along the path", covers a multitude of options for control and intervention. One could unpack that step, but I will argue that at one level, it is not necessary to do so.

From the point of view of control point analysis, either the step succeeds or fails. What is interesting is how to recover from the failure, should it occur. As I argued above, technology cannot ensure success if one is depending on an untrustworthy actor. One must constrain that actor to behave in a minimally acceptable way, or avoid it all together. That is what the technical design must permit, and the details of all the different ways that a bad actor can act badly don't really matter, except to the extent that we try to find that absolute minimum of acceptable behavior that allows the operation to succeed. The diagram contains a good example of that. Imagine that the user is attached to an access ISP (perhaps a untrustworthy hot-spot) that seems to be disrupting communication. If the user opens a VPN, which encrypts all the traffic and sends it back to a trusted point to be decrypted and then sent on, the residual dependency we make on the ISP is that it forward a stream of encrypted (and thus undifferentiated) packets to a destination. An ISP could, of course, block encrypted packets, or block packets to known VPN end-points, and at that point, any further attempts to make use of this ISP become convoluted at best. But this is the space of contention between ISP and user—the detailed analysis of what the ISP can do to attack the flow becomes interesting only in the context of specific mechanisms that attempt to protect the flow and mitigate the intervention. The specific mechanisms prune the options for control, and thus simplify the degree to which the options for control need to be fully analyzed.

Conclusions

The catalog of actors in this discussion has focused on those most directly involved with the creation and operation of cyberspace—the ISPs, the designers of applications, services and content, the users and so on. This discussion has only superficially addressed the actors that sit a bit removed, but which have substantial concerns about the shape of cyberspace, most obviously governments, and as well large industries that find themselves being influenced or reshaped by the cyberspace experience. The examples above about state and private sector influence over the DNS illustrate this influence.

If the actual behavior of the Internet is governed by loosely defined norms, how and where are they enforced? In fact, the operational governance of the Internet is largely located in informal social networks of ISP staff, who meet to get to know (and trust) each other so that when disruptions occur, they can collectively determine what to do to resolve the issue.

Could a different technical design help? Some aspects of the Internet seem fundamental; it is hard (though not impossible) to imagine an Internet without a routing algorithm. But given all the issues that surround the DNS, could an alternative version of the DNS have been designed that would remove some of these options for control all together? Probably yes. There are applications today (including those that are most attentive to the potential of control and adverse influence, such as peer-to-peer music sharing systems) that try to minimize their use of the DNS. The DNS is not a fundamental part of the Internet architecture, just a widely used convenience.

For engineers who design technical mechanisms, it is tempting to look for a purely technical solution that can mitigate some of the undesirable points of control. However, in the presence of untrustworthy actors, this may in general be hard. End-to-end checks between trustworthy endpoints may only be able to confirm reliably that an actor with control has caused a failure. An extreme point of view is that in principle no actor should be considered trustworthy, but this view is not consistent with the way society works, and may make progress impossibly difficult. This reality suggests that the better approach for system design is to incorporate both technical features and the ability to select among actors in order to choose those that are trustworthy[6]. But the design approach that marries technical features with selection of trustworthy actors is not a commonly recognized engineering approach. One conclusion from this analysis is that the approach might deserve more attention.

[6] For an extensive discussion of the role of trust, see Clark, D. D. and M. S. Blumenthal (2011). "The End-to-End Argument and Application Design: The Role of Trust." Federal Communications Law Review **32**(2).

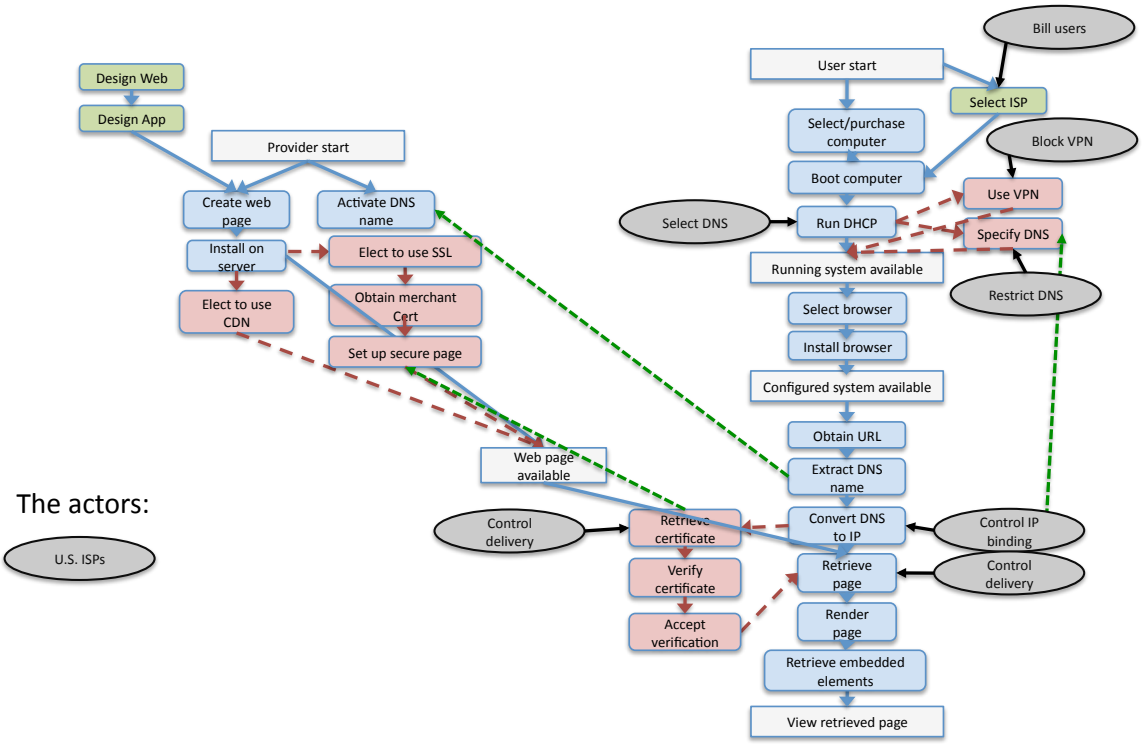

Figure 5: Points of control for a typical ISP in the U.S.

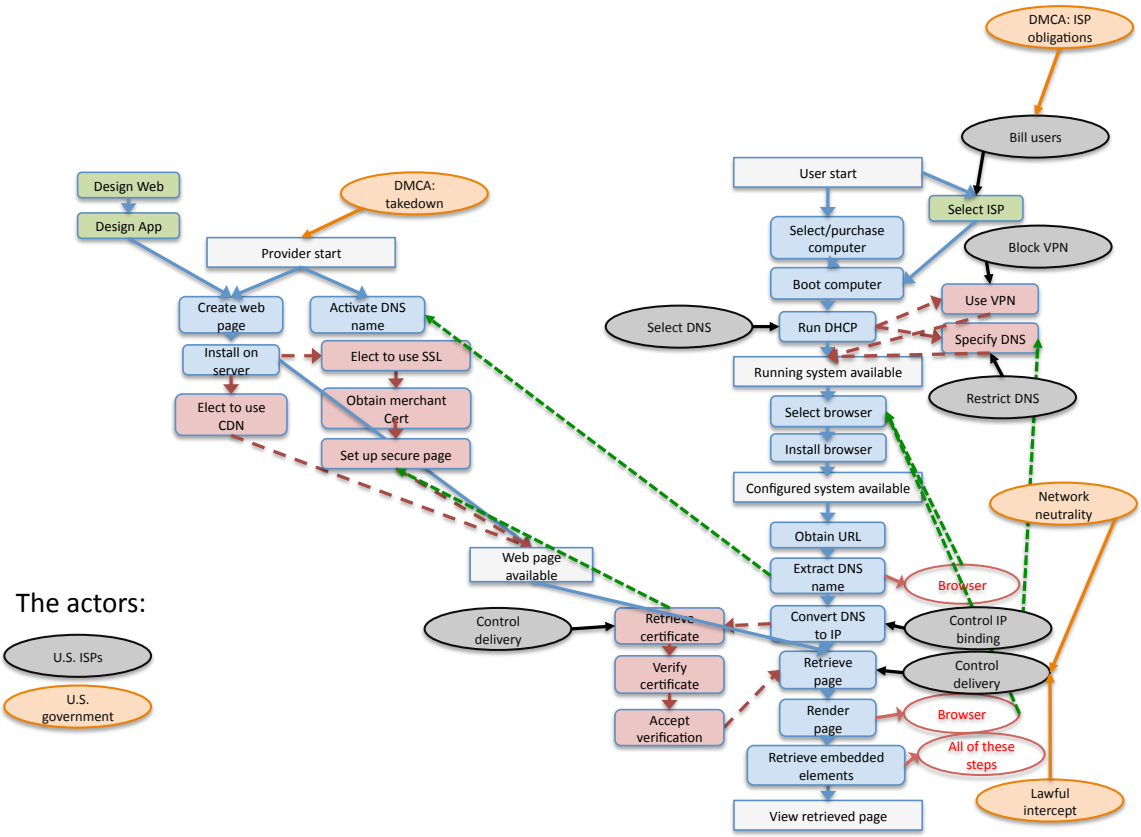

Figure 6: Points of control for the U.S. government, showing influence over ISPs and other primary actors.

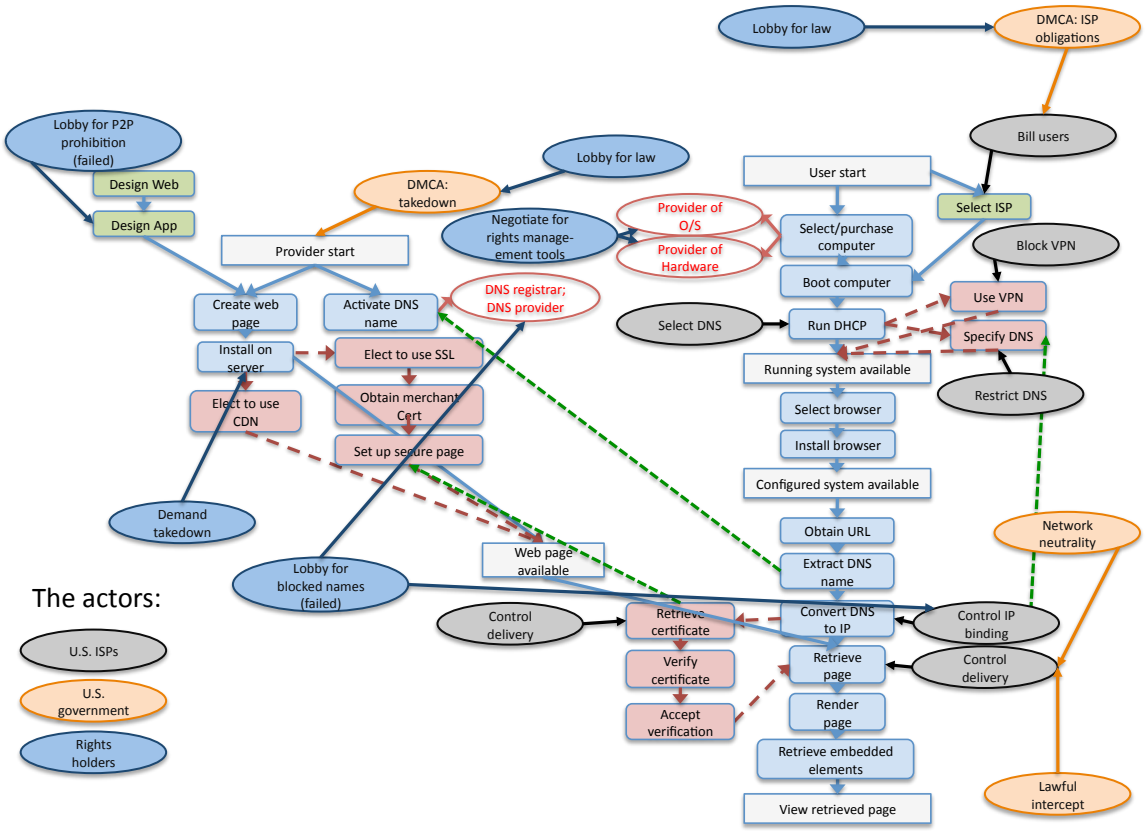

Figure 7: Points of control for a content owner trying to suppress piracy.

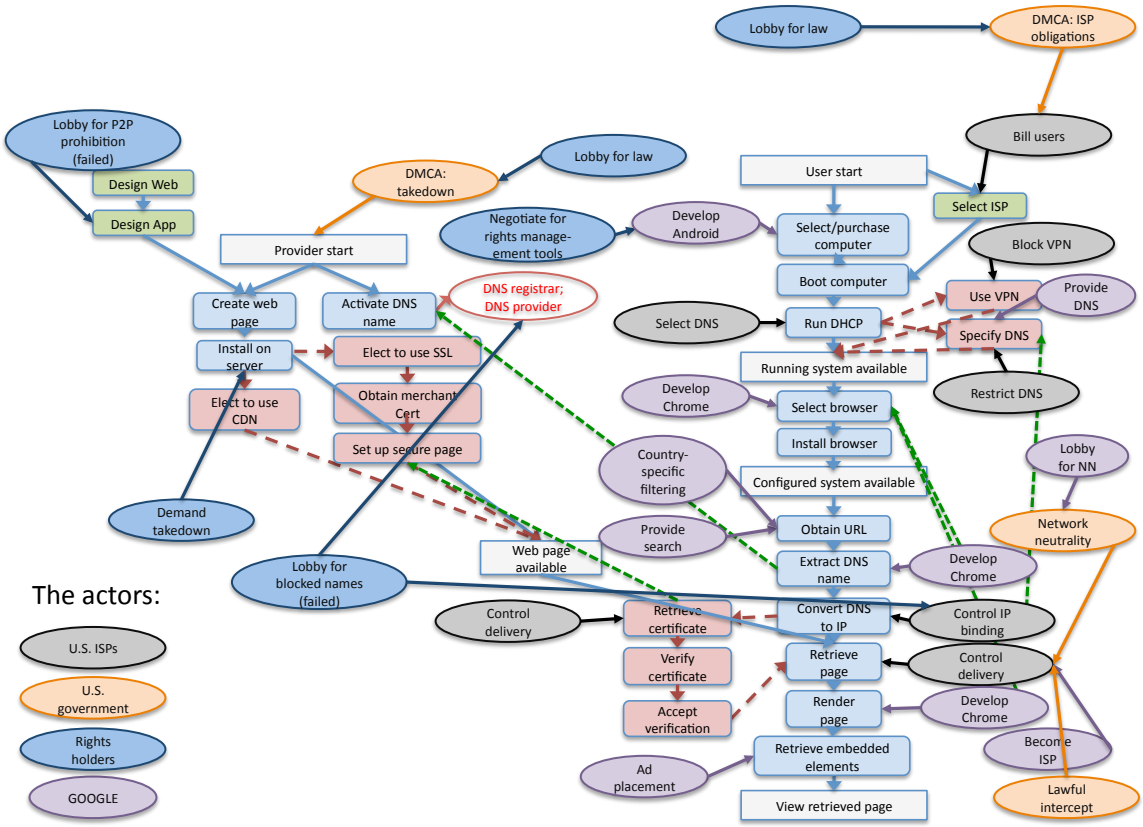

Figure 8: Points of control for Google. Google is a distinctively powerful private sector actor, with many direct means to control the experience of using the Internet.

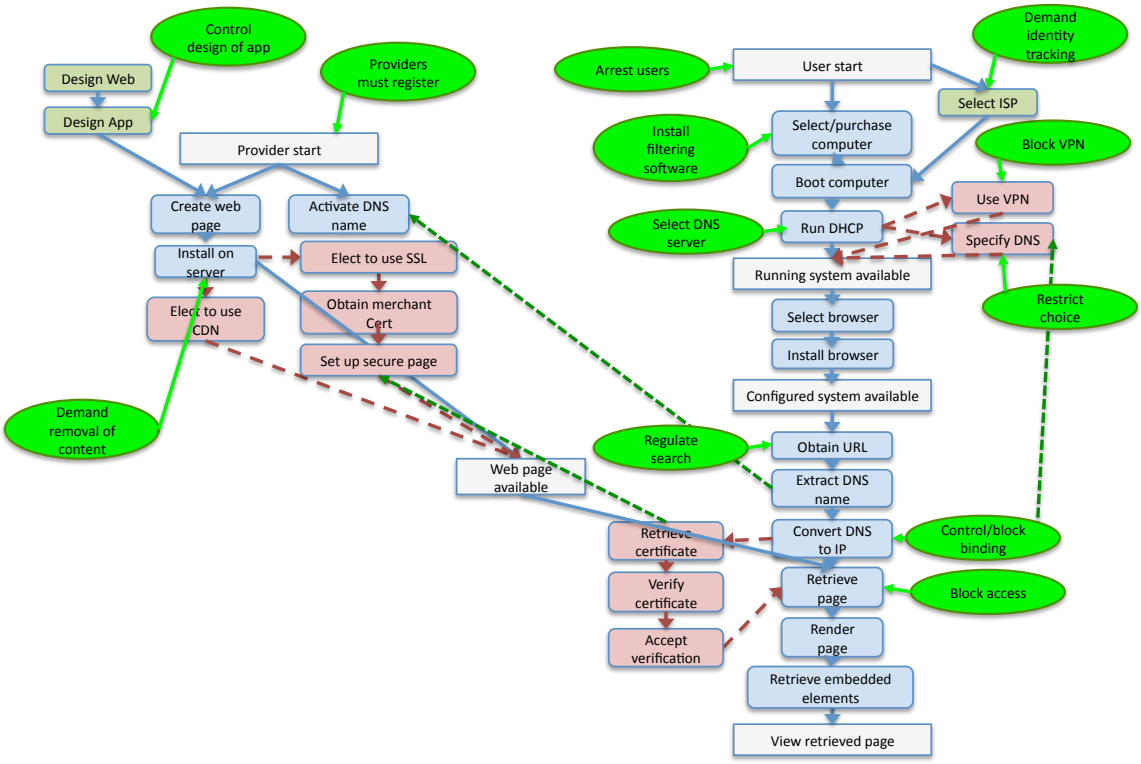

Figure 9: Points of control as used by the government of China. Distinction between primary actor (e.g. the ISP) and the state is not illustrated.