The Great Firewall Revealed

Global Internet Freedom Consortium
http://internetfreedom.org
December 2002

1. Introduction

The past a few months have witnessed the escalating scale of Internet blocking by China, highlighted by the dramatic scenario Google had experienced [1, 2], and by the less fortunate fate Altavista has been suffering [3]. Internet users in China are experiencing inaccessibility to an increasing number of overseas websites which either contain explicit information the authorities do not want the users to see (e.g., www.amnesty.org), or are accidentally victimized by the China's expanding blocking effort (e.g., www.mit.edu).

Recently there have been a few studies by separate groups or individuals to document the current status of China's Internet blocking activity [4, 5, 6]. Zittrain & Edelman [4] have been extensively collecting the websites blocked by China, and their studies reveal a surprisingly large number of sites, most of which are not pornography--related, are inaccessible from within China. A similar study by Villeneuve [6] showed a large portion of the websites related to Falun Gong, Human Rights, Taiwan and Tibet are being blocked.

These efforts took an empirical approach to enumerate the blocking status. This approach provides valuable first-hand observation on the blocking phenomena. However, it can not provide any technical insight to reveal the underlying blocking mechanism. Moreover, these enumeration studies can not give any clue as the websites are blocked by one mechanism or several. To evaluate, understand, circumvent, penetrate or even defeat the Great Firewall, much more detailed technical information is needed, especially in the area regarding the inner working of the Great Firewall.

In this report we document our technical studies on the Internet blocking mechanisms implemented by China. Our expertise in computer networking, security, and programming, our connections with China, and our direct access to test computers inside China, enable us to carry out systematic technical studies of the Great Firewall in unprecedented level of details and depth. The results of our research provided us with solid knowledge of the sophistication, coordination as well as the weakness of the invisible Great Firewall.

Our extensive studies and testing of the Internet blocking infrastructure in China revealed that the blocked websites documented in other studies are actually made inaccessible by three distinct blocking/filtering mechanisms: *IP blocking*, *TCP connection cut-off*, and *URL hijacking*. These three mechanisms are highly effective, and are operating independently, but with certain degree of coordination and some overlap in the functionalities for achieving the overall goal of preventing users from seeing certain websites or content.

In Section 2 we report the technical details of the *IP blocking* mechanism. Section 3 deals with the second mechanism, *TCP connection cut-off.* We demonstrate the *URL hijacking* approach in Section 4. The results are summarized in Section 5. In addition, Appendix A documents the keywords used by the second and third blocking technique to block about 1000 URLs, which are listed in Appendix B.

There are several major Internet operators in China, and in this report our primary focus is CHINANET, which is the largest Internet operator. We believe CHINANET is leading the implementation of Internet blocking technologies, and actively sharing such technologies with other Internet operators. Nevertheless, the other Internet operators are also under active investigation.

2. IP blocking

All the Internet traffic is carried by IP packets, each of which contains an IP number indicating the source address and another IP number indicating the destination address, in addition to its payload and some control data. On its Internet path from the source IP address to the destination IP address, an IP packet goes through a series of routers, which inspect the packet's destination IP address and forward the packet to its next hop accordingly. If a router drops the IP packets destined to certain IP addresses, all the computers networked through this router will not be able to communicate with the machines having these IP addresses, regardless of the content to be transferred by the IP packets. This is essentially how the *IP blocking* works, and we have identified this is one of the three blocking mechanisms at work, which is responsible for a large number of inaccessible websites listed by [4] and [6].

We tested this mechanism using test computers inside China, and use the standard networking utilities such as ping and traceroute (tracert on Windows) to trace the paths an IP packet would take from within China to various overseas websites. Our tests showed that, an IP packet destined to various blocked website would first go through a few area routers, then would be forwarded by some of these routers, and it would eventually forwarded to the following international-level router,

```
p-0-0-0-r1-I-bjbj-1.cn.net (202.97.33.2)
```

at which all the IP packets would be dropped. This observation agrees with the results in [6]. On the other hand, for unblocked websites, an IP packet will take various paths, depending on the desitination, first going through a few area-level routers, then one or two national-level routers, one international-level router, then hopping into the overseas network, and eventually arriving at these websites. These unblocked IP packets will not have a chance to go through the blocking router (`202.97.33.2`) at all.

Following is a sample test case illustrating how the search engine www.altavista.com was blocked by this mechanism. The test was carried out on a computer inside China. The test computer could not "ping" the website, and "tracert" indicated that the IP packets were dropped by the router (`202.97.33.2`).

```
C:\>ping www.altavista.com

Pinging www.altavista.com [209.73.164.90] with 32 bytes of data:

Request timed out.
Request timed out.
Request timed out.
Request timed out.

Ping statistics for 209.73.164.90:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum =  0ms, Average =  0ms

C:\>tracert www.altavista.com

Tracing route to www.altavista.com [209.73.164.90]
over a maximum of 30 hops:

  1    10 ms    10 ms    10 ms  61.179.y.y
  2    10 ms    10 ms    10 ms  10.254.126.1
  3     *      <10 ms    10 ms  10.254.124.5
  4     *        *        *     Request timed out.
  5     *        *        *     Request timed out.
  6    10 ms    10 ms    10 ms  61.179.255.49
  7    20 ms    20 ms    10 ms  202.102.129.253
  8    20 ms    20 ms    30 ms  p-0-0-r1-c-shsh-1.cn.net [202.97.39.1]
  9    81 ms    40 ms    40 ms  p-4-0-r2-c-bjbj-1.cn.net [202.97.34.33]
 10    41 ms    40 ms    50 ms  p-12-0-r1-c-bjbj-1.cn.net [202.97.37.1]
 11    40 ms    40 ms    40 ms  p-1-0-0-r1-i-bjbj-1.cn.net [202.97.33.2]
 12     *        *        *     Request timed out.
 13     *        *        *     Request timed out.
 14     *        *        *     Request timed out.
 15     *        *        *     Request timed out.
 16     *        *        *     Request timed out.
 17     *        *        *     Request timed out.
 18     *        *        *     Request timed out.
 19     *        *        *     Request timed out.
 20     *        *        *     Request timed out.
```

Example 2-1: How www.altavista.com was blocked by *IP blocking*. Test time: 9/25/2002; Test machine IP: 61.179.x.x. At the time of the test then, the URL of www.altavista.com had not been hijacked yet (see Section 3).

For comparison, Example 2-2 below showed the path an IP packet took to reach an unblocked website from the same test machine. The IP packet was not forwarded to the blocking router (`202.97.33.2`) by othe upstream routers.

```
C:\>tracert -h 15 www.google.com

Tracing route to www.google.com [216.239.33.101]
over a maximum of 15 hops:

  1    10 ms    10 ms    10 ms  61.179.y.y
```

```
 2    10 ms    10 ms    10 ms  10.254.126.1
 3   <10 ms    40 ms    10 ms  10.254.124.5
 4     *        *        *     Request timed out.
 5     *        *        *     Request timed out.
 6    10 ms    20 ms    10 ms  61.179.255.49
 7    10 ms    20 ms    10 ms  202.102.129.253
 8    20 ms    20 ms    30 ms  p-0-1-r1-c-shsh-1.cn.net [202.97.39.5]
 9    20 ms    30 ms    30 ms  202.97.33.90
10   150 ms   150 ms   150 ms  202.97.51.2
11   181 ms   190 ms   190 ms  ibr02-p5-1.paix01.exodus.net [206.79.9.121]
12   180 ms   191 ms   190 ms  bbr01-p6-0.sntc03.exodus.net [209.185.9.241]
13   190 ms   180 ms   180 ms  dcr04-g4-0.sntc03.exodus.net [216.33.153.68]
14   180 ms   190 ms   180 ms  csr01-ve241.sntc03.exodus.net [216.33.153.188]
15   180 ms     *      180 ms  google-exodus.exodus.net [64.68.64.210]
```

Example 2-2: The path an IP packet took from the test machine to reach www.google.com. Test time: 9/25/2002; Test machine IP: 61.179.x.x. At the time of the test then, China stopped blocking www.google.com.

We performed numerous similar tests from various test computers, and Figure 2-1 is a partial summary of the test results. We can see that certain routers (indicated by the red IP addresses/hostnames) treated the IP packets differently, depending on the IP was blocked or not. So we speculate these routers have an ACL (access control list) to determine the fate of the IP packets.

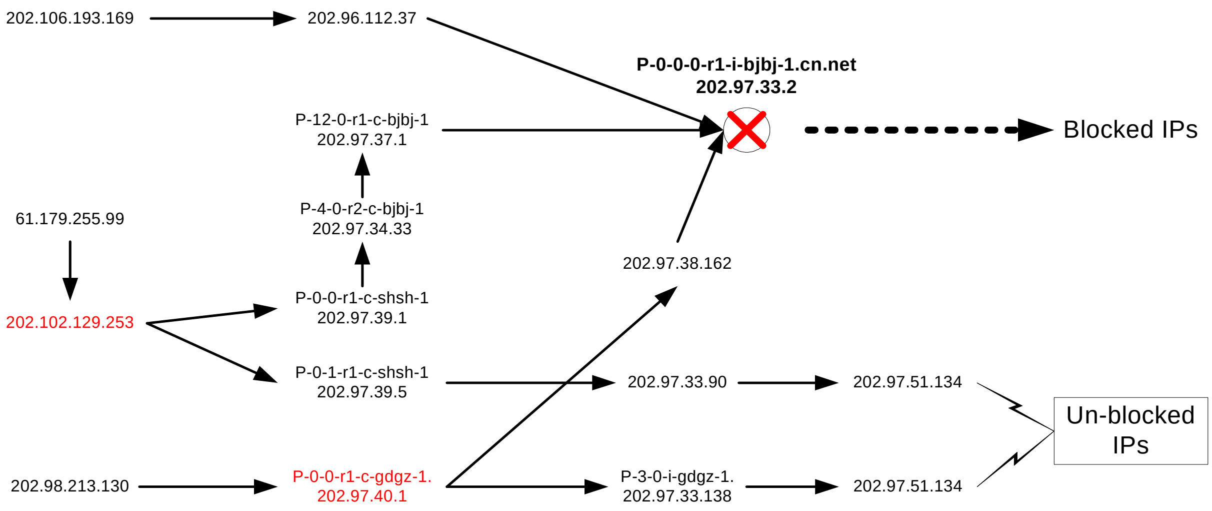

Figure 2-1: The different paths an IP packet took to reach blocked IP addresses and unblocked ones. This plot shows all the IP packets destined for blocked websites were eventually dropped by the router (202.97.33.2).

Our first-hand experiences indicate that *IP blocking* is the earliest method China employed for Internet blocking, and it is still being active, with new IP addresses added to the ACL occasionally. This technique is not substantially technical, but requires human inspection of websites' content to determine if their IP addresses should be added to the ACL or not.

3. TCP connection cut-off

The second blocking mechanism we discovered is *TCP connection cut-off.* When a user visits a website, his/her computer will first establish a logical data connection using the TCP protocol, which is built upon the IP packets. The set-up of a TCP connection between two computers involves a three-way hand-shaking process, and once this process is successfully completed, information will be wrapped in TCP segments and will flow through this connection in either direction. These TCP segments, like the IP packets, will also go through a series of routers on their path from the source to the destination. A router on such a path can cut off such a TCP connection by sending to either end a TCP segment containing a "RESET" flag, consequently blocking the information flow between the two computers in question. This is exactly what China is using as another road block on the information highway.

We used numerous tools for this study, including standard TCP connection testing utilities such as netcat, telnet; network protocol analyzers such as Ethereal; and custom-developed raw packet generating and sniffing programs.

Our studies revealed that this blocking mechanism is consisted of the following components:

- Sniffing of TCP segments
- Pattern matching between TCP data and a keyword list
- Instant reset
- Stateful reset

A router performing such a blocking task inspects (sniffing) each TCP segments passing through, and compares the data contained in each TCP segments with a list of pre-defined keywords (pattern matching). If the router spots a match, it will generate two TCP RESET segments, destined to the source and destination, respectively, to interrupt the TCP connection (instant reset) if one does exist. Furthermore, if the offending TCP segment is sent in the context of an established TCP connection, this TCP connection will be remembered by this router in terms of a triad (source IP, destination IP, destination port). Any subsequent attempts to re-establish a TCP connection involving this triad will be re-set (stateful reset), right after the three-way hand-shaking process is completed. Such memory lasts approximately **150** seconds, after which TCP connections with the same triad can be established successfully and data can be transmitted as long as not keywords are matched. In case a match is spotted, the blocking kicks in again.

The TCP RESET segments generated by these blocking routers have distinct fingerprints, which made them easy to be recognized from other normal RESET segments. These segments have a unusual window size (ws =1), and the TTL of the underlying IP packets has a extremely low value of 40 to 55.

Our tests show that this blocking mechanism is not limited to the web traffic. The router inspects ANY TCP segments. So it is possible that if email traffic matches one of the keywords, the connection will be cut off too.

Table 3-1 shows a summary of our test results with both the instant reset and stateful reset, for both HTTP traffic (port 80) and non-HTTP traffic. We have identified "tibet.org" is one of keywords, and we either send a stand-alone TCP segment to a test machine inside China, to test the instant reset, or, first establish a TCP connection with the test machine, then send a TCP segment with the sensitive content, to test the stateful reset. These tests showed that the *TCP cut-off* mechanism does not differentiate HTTP or non-HTTP traffic. It sniffs ANY TCP traffic. Both the instant reset and stateful reset recognizes the same keyword pattern, which can be expressed in a standard Unix regular expression as follows,

```
/^GET \/.*keyword.*/i
```

Table 3-1: Test of keyword pattern matching for both instant reset and stateful reset, for both HTTP and non-HTTP traffic.

Destination IP:port	**202.108.249.206:80**		**202.108.44.208:110**	
keywords	Instant reset	Stateful reset	Instant reset	Stateful reset
"GET /tibet.org"	**Y**	**Y**	**Y**	**Y**
" GET /tibet.org"	N	N	N	N
"GET /tibet.org"	N	N	N	N
"get /tibet.org"	**Y**	**Y**	**Y**	**Y**
"GET tibet.org"	N	N	N	N
"tibet.org"	N	N	N	N
"GET /tibet"	N	N	N	N
"GET /tibet.orgCCCC"	**Y**	**Y**	**Y**	**Y**
"GET /CCCCCtibet.org"	**Y**	**Y**	**Y**	**Y**
"GET /CCCCCtibet.orgDD"	**Y**	**Y**	**Y**	**Y**
"GET tibet.orgCCCC"	N	N	N	N
"HEAD /tibet.org"	N	N	N	N
"minghui"	N	N	N	N
"GET /minghui"	**Y**	**Y**	**Y**	**Y**
"minghui"(Chinese GB)	N	untested	N	untested
"Falun Gong"(Chinese GB)	N	untested	N	untested

Following is a list of sample keywords. For a comprehensive list, see Appendix A, which contains more than 250 keywords as we identified at the time of this report.

```
altavista.com
amnesty.org
bignews
boxun.com
cnd.org
epochtimes.com
```

```
falun
fawanghuihui.org
fgmtv.org
hongkong.com
minghui
mit.edu
myftp.com
myftp.org
no-ip
renminbao.com
secretchina.com
taiwan.com
tibet.com
tibet.net
tibet.org
voa.gov
vot.org
wenxuecity.com
xinsheng.net
zhengjian.org
zhengwunet.org
```

List 3-1: A sample of the keywords used by the *TCP cut-off* mechanism. The same set of keywords is shared by the *URL hijacking* mechanism (see Section 4). A comprehensive list of the keywords is in Appendix A.

Based on our knowledge of CHINANET's network topology, we were also able to map out those routers which were participating the *TCP cut-off* operation, as shown in Figure 3-1 below. We can see this mechanism is implemented as a distributed system, with 10 CHINANET national-level routers geographically located in BeJing, ShangHai and GuangZhou. The routers' IP addresses have the form 202.97.33.oddnumber, possibly indicating centralized management.

4. URL hijacking

The third blocking mechanism, *URL hijacking*, is what China is using to disrupt DNS (Domain Name Service) in order to hijack a website's URL and redirect visitors to a different, inaccessible destination. Normally when a user types in a URL in his/her browser, the user's computer will need first to convert the URL (such as www.mit.edu) to a numerical IP address (which is 18.181.0.31) in order to send IP packets to it and carry out TCP-based communication. To do so, the computer will issue a DNS type "A" query to its pre-defined DNS resolver, and the DNS resolver, if it does not know the answer, will in turn issue a series of DNS queries to locate the domain mit.edu's authoritative DNS server, and eventually the user's resolver will issue a type "A" query to mit.edu's authoritative DNS server, which is located outside China. Such a query is contained in a UDP datagram, which is a special form of IP packet, and the payload contains the URL "www.mit.edu" and an flag indicating the type "A" query. Upon receiving such a query, the mit.edu's authoritative server will respond with the correct IP address of the URL "www.mit.edu" .

China’s URL hijacking mechanism is implemented in such a way that a number of the CHINANET routers constantly sniffing the passing network traffic and look for keywords in DNS type “A” queries. Once they see such a query containing a URL which matches one of the keywords in the list, these routers will issue a fake reply, disguised as the reply from the authoritative DNS server (spoofing). Since the offending DNS query is not dropped by the routers, it will eventually reach the authoritative server, and the server will issue a true answer. So the user’s computer will actually receive two DNS answers if it issues a query with a keyword in it, and the computer always take the first answer, which is apparently the fake one, and subsequently the user’s browser is directed to a wrong location dictated by the fake reply.

For the hijacking routers in CHINANET, the spoofed fake replies are always the same IP address:

64.33.88.161

which is an overseas IP address and already blocked by *IP blocking* (Section 1). The end result for the user is a error page generated by the browser complaining the page can not be found. The hijacking routers do not care if the offending query is coming from inside China or the opposite – they operation bi-directionally.

Following is an example demonstrating how DNS queries containing the keyword “no-ip” were hijacked and answered with fake IP addresses. The keyword matching pattern can be expressed as a Unix regular expression:

/.***keyword**.*/i

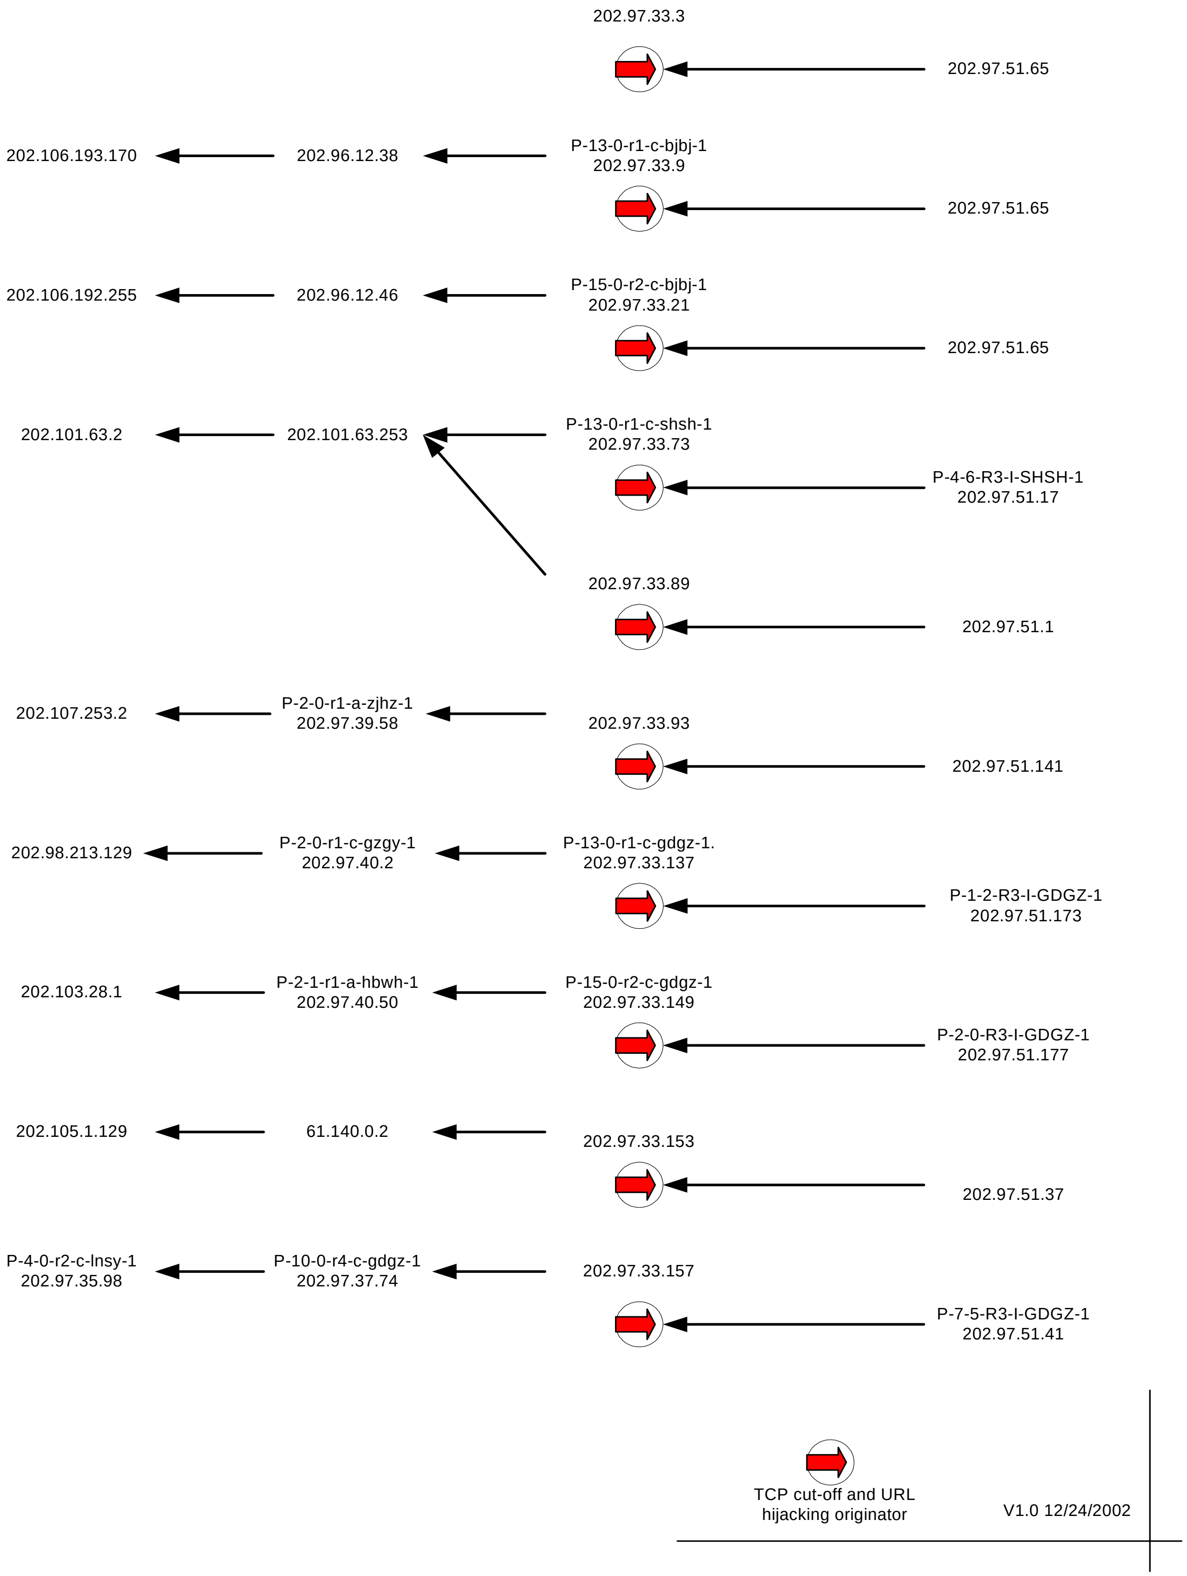

Figure 3-1: CHINANET routers which perform TCP cut-off and URL hijacking (see Section 4) functions. There routers are both logically and geographically distributed, sharing the same list of keywords, however.

```
> server 202.106.186.229
Default Server:  [202.106.186.229]
Address:  202.106.186.229

> no-ip
Server:  [202.106.186.229]
Address:  202.106.186.229

Non-authoritative answer:
Name:    no-ip.localdomain
Address:  64.33.88.161

> no-IP.com
Server:  [202.106.186.229]
Address:  202.106.186.229

Non-authoritative answer:
Name:    no-IP.com
Address:  64.33.88.161

> NO-IP.org
Server:  [202.106.186.229]
Address:  202.106.186.229

Non-authoritative answer:
Name:    NO-IP.org
Address:  64.33.88.161

> AAAAAno-ipCCC
Server:  [202.106.186.229]
Address:  202.106.186.229

Non-authoritative answer:
Name:    AAAAAno-ipCCC.localdomain
Address:  64.33.88.161

> AAA1.no-ipBBB.CCCC
Server:  [202.106.186.229]
Address:  202.106.186.229

Non-authoritative answer:
Name:    AAA1.no-ipBBB.CCCC
Address:  64.33.88.161
```

Example 4-1: How DNS type “A” queries containing the keyword “no-ip” were hijacked with fake replies. The test was performed from an overseas computer using the standard “nslookup” program, and using a DNS resolver inside China, 202.106.186.229.

With in-depths tests we discovered that the *URL hijacking mechanism* and the *TCP cut-off* mechanism (Section 3) are actually implemented on the same group of CHINANET routers. In other words, a router words, a router doing *URL hijacking* also performs *TCP cut-off*, and *vice versa*, as shown in Figure 3-1. Further more, the keyword list is the same for both mechanisms. So any URL containing any keyword in List 3-1 will be hijacked. This has some surprising consequences. For example, since “hongkong.com” is a keyword in the list, all URLs containing this keyword will be hijacked, such as the following URLs, some of which, we believe, are not their intention to block,

```
car.hongkong.com
cityguide.hongkong.com
```

```
home4u.hongkong.com
hongkong.com
lifestyle.hongkong.com
shop4u.hongkong.com
wap4u.hongkong.com
webserv1.discoverhongkong.com
women.hongkong.com
www.autism-hongkong.com
www.bighongkong.com
www.childhealthhongkong.com
www.discoverhongkong.com
www.explore-hongkong.com
www.food4hongkong.com
www.helihongkong.com
www.home4u.hongkong.com
www.hongkong.com
www.hotelshongkong.com
www.regalhongkong.com
www.stanfordhongkong.com
```

List 4-1: URLs being hijacked due to the presence of the keyword "hongkong.com", regardless the presence of this keyword is by design or purely by accident**.**

Appendix A contains a comprehensive list of the keywords, and a long list of hijacked URLs, which contain or happen to contain one of these keywords, can be found in Appendix B.

5. Summary

A number of recent empirical studies on China's Internet blocking have identified a surprisingly large number of websites being blocked. However, these results do not provide any information as how the blocking works. This report documented, in great technical detail and depth, the mechanisms China has implemented to block un-wanted Internet traffic. We discovered that the offending websites were actually blocked by three different but related mechanisms, *IP blocking*, *TCP connection cut-off*, and *URL hijacking*. We covered all the technical aspects of these mechanism, and significantly escalate our knowledge of the state-of-the-art blocking technology.

Such knowledge is fundamental to our ability to keep track of the ever-evolving blocking technology China is deploying, and is crucial in designing techniques to circumvent, penetrate, or defeat the "Great Firewall". For example, our studies showed that since the URL hijacking is bi-directional, a user inside China still can not get the correct IP of a URL even he/she uses a DNS server outside China, in contrary to the suggestion made in [4]. As a matter of fact, based on our unprecedented technical understanding of the blocking mechanisms, we have designed various counter-measures for web browsing, emailing, etc, and these counter-measures have been approved highly successful. These results will be documented in a separate report.

References:

[1] BBC News, 2002: “China blocking Google”.
HTTP://news.bbc.co.uk/2/hi/technology/2231101.stm

[2]Fox News, 2002: “China Ends Blocking of Google Search Engine”.
http://www.foxnews.com/story/0,2933,62859,00.html

[3] CNN, 2002: **“**China blocks AltaVista search engine”.
HTTP://www.cnn.com/2002/TECH/internet/09/06/china.google/

[4] Zittrain & Edelman, 2002: “Empirical Analysis of Internet Filtering in China”.
HTTP://cyber.law.harvard.edu/filtering/china/

[5] Dynamic Internet Technology, Inc., 2002: “Forbidden sites hijacked all over China”.
HTTP://www.dit-inc.us/report/hj.htm

[6] Villeneuve, 2002: “Project C (r. 1.0)”.
HTTP://www.chass.utoronto.ca/~citizenl/assets/articles/ProjectC-r1.pdf