ISP Content Filtering: methods, failures and some politics

Richard Clayton

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Leuven

15th November 2006

Summary

- Content blocking system taxonomy
- Overblocking and avoidance
- Cleanfeed and the “oracle attack”
- The Great Firewall of China
- The political landscape

Taxonomy (blocking methods)

- DNS poisoning
  - refuse to resolve the wicked domains
  - low cost, and highly scalable
- Blackhole routeing
  - refuse to carry the traffic to the wicked site
  - low cost, but limits to size of ACLs/routing-table
- Proxy filtering
  - refuse to serve the wicked pages
  - high cost, and all traffic has to be inspected

Problems with DNS poisoning

- Apparently easy…

```
@ IN SOA localhost. root.localhost. (
        2004010100 86400 3600 604800 3600 )
@ IN NS  localhost.
@ IN A   127.0.0.1
* IN A   127.0.0.1
```

- But getting it right for subdomains and for email requires some thought! Dornseif found that every German ISP he studied had made errors!

Problems with blackhole routeing

- Dropping packets will (obviously) affect every website hosted at the IP address!
  - hence useless for geocities.com
  - in fact useless for huge numbers of other sites as well. Edelman study found “overblocking” a significant issue: 87.3% of com/net/org sites share IP address with at least one other; 69.9% with at least 50 others (and a continuum exists at all sizes)
  - do you really want to block the “Romanian Tourist Board” website ?

Problems with proxy filtering

- This method avoids overblocking (huzzah!)
- However, it can have significant costs in equipment, in customer satisfaction and in network reliability
  - economic justifications for caching proxies continue to get weaker
  - proxies often slower than going direct!
  - caching proxies obstruct many personalisation schemes for website content providers

Avoidance for clients

- Use a different DNS server
- Use IP addresses
- Use a relay (often encrypts and anonymises)
- Encode request%73 to avoid recognition
  - look at your spam to see this raised to an art form
- Send malformed HTTP requests
  - eg: multiple HOST protocol elements

Avoidance for servers

- Move site to another IP address (easy)
- Change port number (hard to discover)
- Provide same content on many different URLs
- Accept unusually formatted requests
  - servlets at client could obfuscate or encrypt so that an intermediary has no chance of using anything short of the IP address to identify content

The IWF

- Internet Watch Foundation
- Set up 1996 in the UK to address problem of child pornography on Usenet
- Operates a consumer “hot-line” for reports
- Now mainly concerned with websites
- Has a database of sites not yet removed
- Database underpins blocking system

Design of CleanFeed

- Part of BT “anti-child-abuse initiative”
  - two stage (hybrid) system, BT, June 2004
- First stage is IP address based
  - candidate traffic for blocking is redirected
- Second stage matches URLs
  - redirected traffic passes through a web proxy
- Best of both worlds?
  - highly accurate
  - but can be low cost because #2 is low volume

Design of CleanFeed

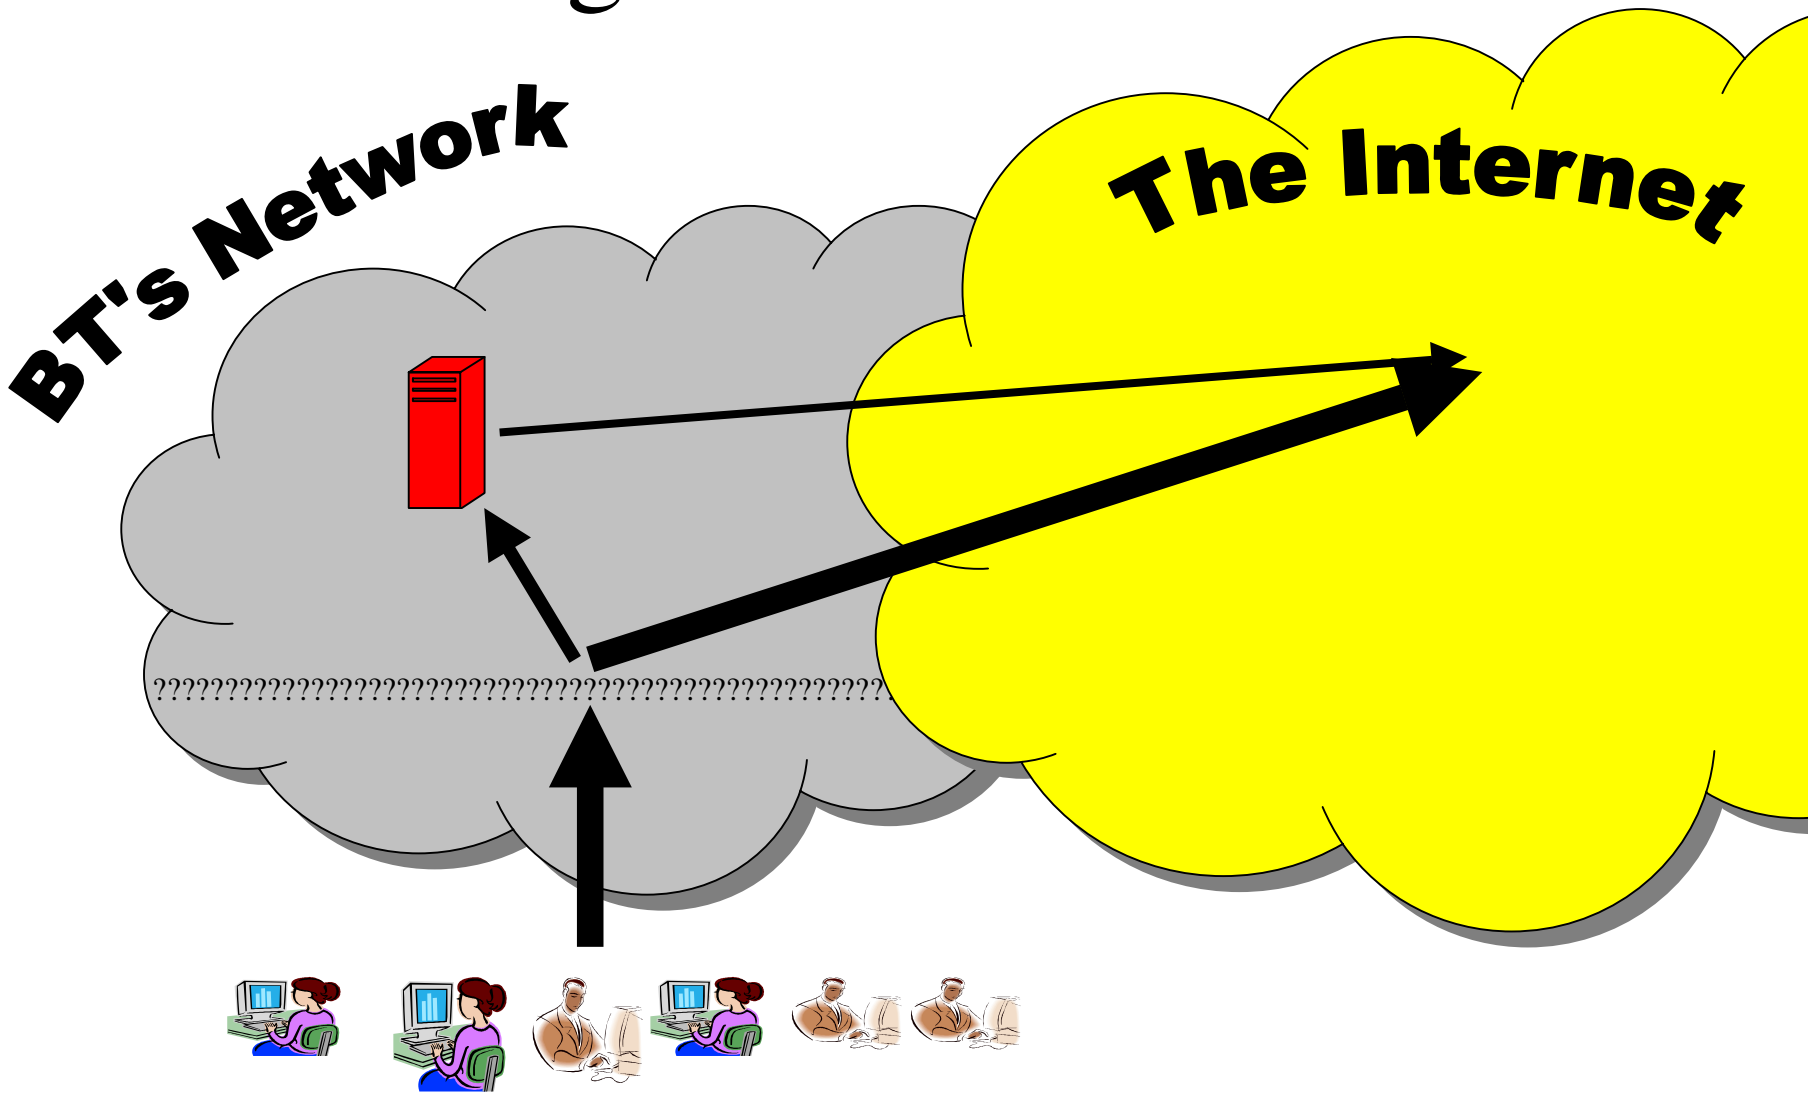

So it’s an elegant design…

… are there any problems with it ?

YES!

Can attack the system

- Redirect extra traffic
  - add specious IP addresses into DNS lookup so that high bandwidth sites are sent to stage #2
- Block valid traffic

  google cache: **`66.102.9.104/search=?q=cache:FF9etc`**

  ‘etc venues’: **`195.224.53.128/directions/parkstreet`**
- NB: more efficient when sure is the IWF

Detecting IWF accesses

- Content providers can self-report
  - provides valuable info about timing etc
  - NB: recognising CleanFeed also relevant
- IWF have a fixed /26 network
  - need anonymising systems (caches, Tor, JAP..)
- Detect multiple accesses for same identifier
  - first AS is (outraged) consumer, second IWF, third the police or other investigators

The oracle attack

- Detect the redirection by the first stage by seeing what traffic reaches the second
- Send `tcp/80` packets with TTL set to 8, see what then comes back:

The oracle attack

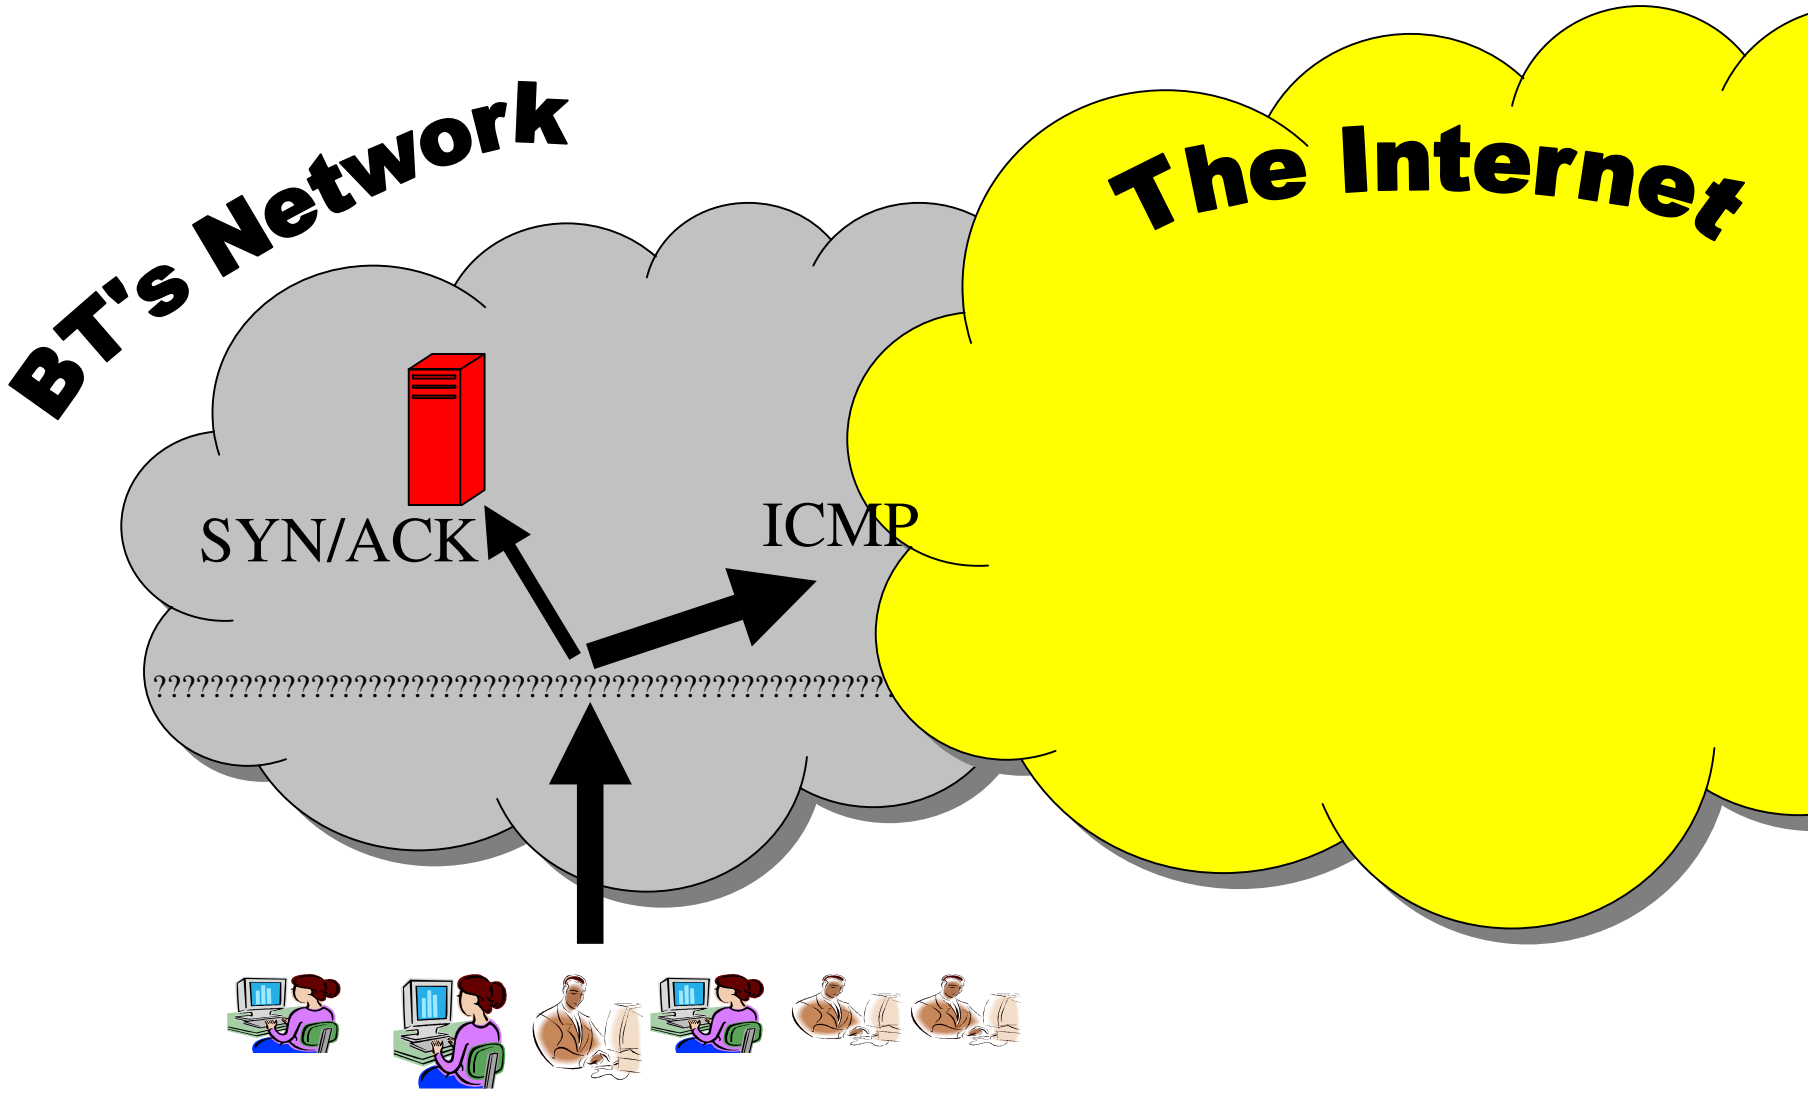

The oracle attack

- Detect the redirection by the first stage by seeing what traffic reaches the second
- Send **`tcp/80`** packets with TTL set to 8, see what then comes back:
  - ICMP time exceeded means no redirect
  - RST (or SYN ACK) means redirect to proxy
- Then use a suitable database to get domain names, eg: **`whois.webhosting.info`**

Oracle attack results I

```
17:54:28  Scan: To [~~~.~~~.191.38] : [166.49.168.9],  ICMP
17:54:28  Scan: To [~~~.~~~.191.39] : [166.49.168.1],  ICMP
17:54:28  Scan: To [~~~.~~~.191.40] : [~~~.~~~.191.40], SYN/ACK
17:54:28  Scan: To [~~~.~~~.191.41] : [166.49.168.13], ICMP
17:54:28  Scan: To [~~~.~~~.191.42] : [~~~.~~~.191.42], SYN/ACK
17:54:28  Scan: To [~~~.~~~.191.43] : [166.49.168.9],  ICMP
17:54:28  Scan: To [~~~.~~~.191.44] : [166.49.168.5],  ICMP
17:54:28  Scan: To [~~~.~~~.191.45] : [166.49.168.9],  ICMP
17:54:28  Scan: To [~~~.~~~.191.46] : [166.49.168.13], ICMP
17:54:28  Scan: To [~~~.~~~.191.47] : [166.49.168.9],  ICMP
17:54:28  Scan: To [~~~.~~~.191.48] : [166.49.168.9],  ICMP
17:54:28  Scan: To [~~~.~~~.191.49] : [~~~.~~~.191.49], SYN/ACK
17:54:28  Scan: To [~~~.~~~.191.50] : [~~~.~~~.191.50], SYN/ACK
```

Oracle attack results II

```
~~~.~~~.191.40   lolitaportal.****
~~~.~~~.191.42   no websites recorded in the database
~~~.~~~.191.49   samayhamed.****
~~~.~~~.191.50   amateurs-world.****
                 anime-worlds.****
                 boys-top.****
                 cute-virgins.****
                 cyber-lolita.****
                 egoldeasy.****
                 elite-sex.****
              ... and 26 more sites with similar names
```

NB: missing names probably `.ru` or outdated database

NB: dodgy names on `.41` `.43` … BUT no IWF “endorsement”

NB: It is illegal for me to check the ACTUAL contents

(Not) fixing the oracle attack

- There were other two-stage systems deployed in the UK (unknown to me)
- The oracle attack worked there too!
- Attempted to fix them by discarding all packets with low TTL
- Scanning program rewritten to examine TTL on *incoming* packets instead!
- It is never going to be possible for a nearby proxy to perfectly emulate remote servers!!

The Great Firewall of China

Joint work with Steven Murdoch & Robert Watson

+

assistance was provided for logging etc by a Chinese citizen [who was unaware of what we proposed to do]. Their site does NOT contain any material that should be censored and no censorable requests were made from the Chinese end of the connection.

Keyword filtering

- Chinese firewall shuts connections if it spots specific keywords passing by
  - for example **`GET /?falun HTTP/1.0`**
- Keywords spotted as they pass by in both directions (dealing with requests & results)
- ***CAUTION:*** parts of Chinese system DO use other blocking methods, and the academic network isn't currently using the scheme, and other protocols are blocked at the application level!

Actual mechanism

```
cam(54190) → china(http)[SYN]
china(http)→ cam(54190) [SYN, ACK] TTL=39
cam(54190) → china(http)[ACK]
cam(54190) → china(http) GET /?falun HTTP/1.0<crlf><crlf>
china(http)→ cam(54190) [RST] TTL=47, seq=1, ack=1
china(http)→ cam(54190) [RST] TTL=47, seq=1461, ack=1
china(http)→ cam(54190) [RST] TTL=47, seq=4381, ack=1
china(http)→ cam(54190) HTTP/1.1 200 OK (text/html)<crlf>..
cam(54190) → china(http)[RST] TTL=64, seq=25, ack zeroed
china(http)→ cam(54190) . . . more of the web page
cam(54190) → china(http)[RST] TTL=64, seq=25, ack zeroed
china(http)→ cam(54190) [RST] TTL=47, seq=2921, ack=25
```

Meanwhile…

- The other end of the connection is ***also*** seeing RST packets from the firewall!

Ignoring the firewall

- **Q:** Since the packets pass through the firewall, what happens if the RST packets are ignored?
- **A:** Web page is transferred just fine (though you get a LOT more RSTs as well)
- NB: necessary to ignore RST packets at ***both*** ends of the connection

Further connections

- Trying to connect again causes RST packets to be sent immediately (even if no “bad” keywords are transferred)

```
cam(54191) → china(http)[SYN]
china(http)→ cam(54191) [SYN, ACK] TTL=41
cam(54191) → china(http)[ACK]
china(http)→ cam(54191) [RST] TTL=49, seq=1
```

- Once again dropping RSTs allows transfer

Denial of service attack

- Send single packets (containing `falun`) to Chinese firewall, forging source & destination
- Connection from source to destination blocked
- Single dialup connection can knock many hundreds of connection over
- NB: only pairs of addresses
- NB: only nearby port numbers (? NAT ?)

Firewall design

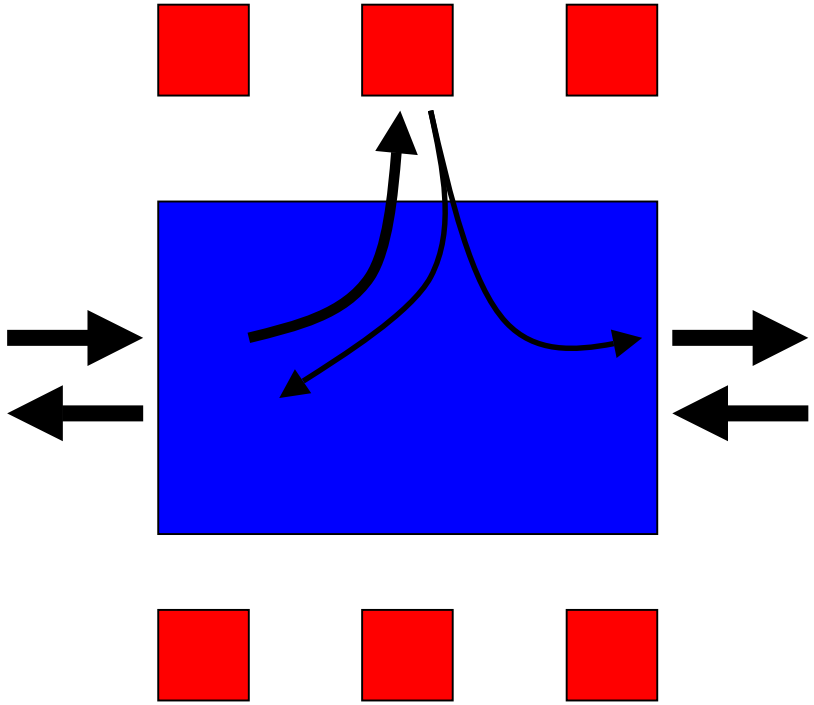

Evidence:

- RST sometimes precedes & sometimes follows data
- RST values (+0, +n, +3n)
- Read the user manuals from (?)providers
- Shuffling of RSTs when a sudden burst of packets

NB: NO STATE IN FIREWALL!

Firewall “state”?

- Splitting `?falun` across packets avoids detection (a surprise! hardware thought to be used can handle this (and overlaps!))
- Refined view is that firewall doesn’t assume it sees packets in both directions, so must do the best it can with the packet in its hand
- Future work will refine our explanation

False SYN/ACKs

```
cam(38104) → china(http)[SYN]
china(http)→ cam(38104) [SYN, ACK] TTL=105
cam(38104) → china(http)[ACK]
cam(38104) → china(http) GET / HTTP/1.0<crlf><crlf>
china(http)→ cam(38104) [RST] TTL=45, seq=1
china(http)→ cam(38104) [RST] TTL=45, seq=1
china(http)→ cam(38104) [SYN, ACK] TTL=37
cam(38104) → china(http)[RST] TTL=64, seq=1
china(http)→ cam(38104) [RST] TTL=49, seq=1
china(http)→ cam(38104) [RST] TTL=45, seq=3770952438
china(http)→ cam(38104) [RST] TTL=45, seq=1
china(http)→ cam(38104) [RST] TTL=45, seq=1
china(http)→ cam(38104) [RST] TTL=45, seq=1
china(http)→ cam(38104) [RST] TTL=45, seq=1
```

Fixing “blocking with confusion”

- Fake SYN/ACK does not confuse once real SYN/ACK has been accepted
- SYN/ACK ***currently*** easy to distinguish
- Real fix is for stack (or a bastion firewall) to hold alternative views of remote sequence value, avoid using a value until see further evidence
  - lack of state in Great Firewall makes this easy(ish)

Porn vs Politics

- Firewall capable of logging events
- No different from encryption/proxies – **but** firewall knows if you're looking at porn or at politics: so may affect your sentence
- Special code is evidence on your machine
- Much better if stack vendors made special tools unnecessary; and there's technical reasons to wish to drop fake resets

Some more general comments….

UK Politics

- Blocking was considered “impossible” until BT deployed CleanFeed
- ISPA claim 80% of consumers covered by systems that block illegal child images
- Minister now wants all of (broadband) industry to be blocking by the end of 2007
  - voluntary except: “*If it appears that we are not going to meet our target through co-operation, we will review the options*”

Whitehall comprehension?

- *"Recently, it has become technically feasible for ISPs to block home users' access to websites irrespective of where in the world they are hosted"*
- In my view, doubtful that they understand the cost, fragility or ease of evasion of these blocking systems, let alone the reverse engineering of the blocking lists.

Other uses?

- Fratini (EU) wants Internet to be a "hostile environment" for terrorists
  - *"I think it's very important to explore further possibilities of blocking websites that incite to commit terrorist action"*
- Drugs, gambling, holocaust denial…
- and don't overlook civil cases:
  - such as, defamation, copyright material, industrial secrets, home addresses of company directors, lists of MI6 agents…

Other countries

- Norway, Sweden & several others blocking child pornography
- Italy blocking gambling sites
- Denmark (Tele2) blocking **`allmymp3.com`**
- Saudi Arabia, Singapore, Burma, and many central Asian countries blocking political speech… see: OpenNetInitiative for info

Conclusions

- Three basic ways of blocking content
- Many (and deep) flaws come from relying on validity of content providers data
- Hybrid systems can be lower cost, but have some extra problems (extracting the site list)
- A key part of the Great Firewall of China relies on acquiescence by the end-points
- Blocking illegal images is top of a very slippery slope, and systems will be used for many things

ISP Content Filtering: methods, failures and some politics

http://www.cl.cam.ac.uk/~rnc1/

PhD Thesis (see Chapter 7) is Tech Report #653

plus two PET Workshop papers, 2005 & 2006

UNIVERSITY OF CAMBRIDGE
Computer Laboratory