ISP Content Filtering: methods, failures and some politics

Dr Richard Clayton

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

LINX55, London

21st November 2006

Summary

- Content blocking system taxonomy
- Avoiding the blocking altogether
- Attacking the blocking systems
- Cleanfeed and the “oracle attack”
- The IWF website list
- The political landscape

Method #1: DNS poisoning

- DNS poisoning
  - refuse to resolve the wicked domains
  - low cost, and highly scalable
- Apparently easy, just add a local SOA in **`bind`**
- But getting it right for subdomains and for email (MX records) requires some thought! Dornseif found that every German ISP he studied had made errors!

Method #2: Blackhole Routeing

- Refuse to carry traffic to/from the wicked site
  - step-cost (limits to size of ACLs/routing-table)
- Affects every website hosted at the IP address!
  - hence useless for geocities.com; and for many other sites as well. Edelman measured "overblocking": 87.3% of com/net/org sites share IP address with at least one other; 69.9% with at least 50 others
- Do you really want to find that you're blocking the "Romanian Tourist Board" website?

Method #3: Proxy Filtering

- Proxy refuses to serve the wicked pages
  - high cost, and all traffic has to be inspected
  - but it does avoid overblocking (huzzah!)
- However, significant costs in equipment, in customer satisfaction and in network reliability
  - economic case for caching proxies ever weaker
  - obstructs many server personalisation schemes
  - proxies often slower than going direct!

Avoidance for clients

- Use a different DNS server
- Use IP addresses rather than hostnames
- Use a relay (often encrypts and anonymizes)
- Encode request%73 to avoid recognition
  - look at email 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 (harder 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

CleanFeed

- Part of BT “anti-child-abuse initiative”
  - two stage (hybrid) system, BT, June 2004
  - similar designs deployed by other ISPs
- 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?
  - accurate, but low cost because #2 is low volume

Design of CleanFeed

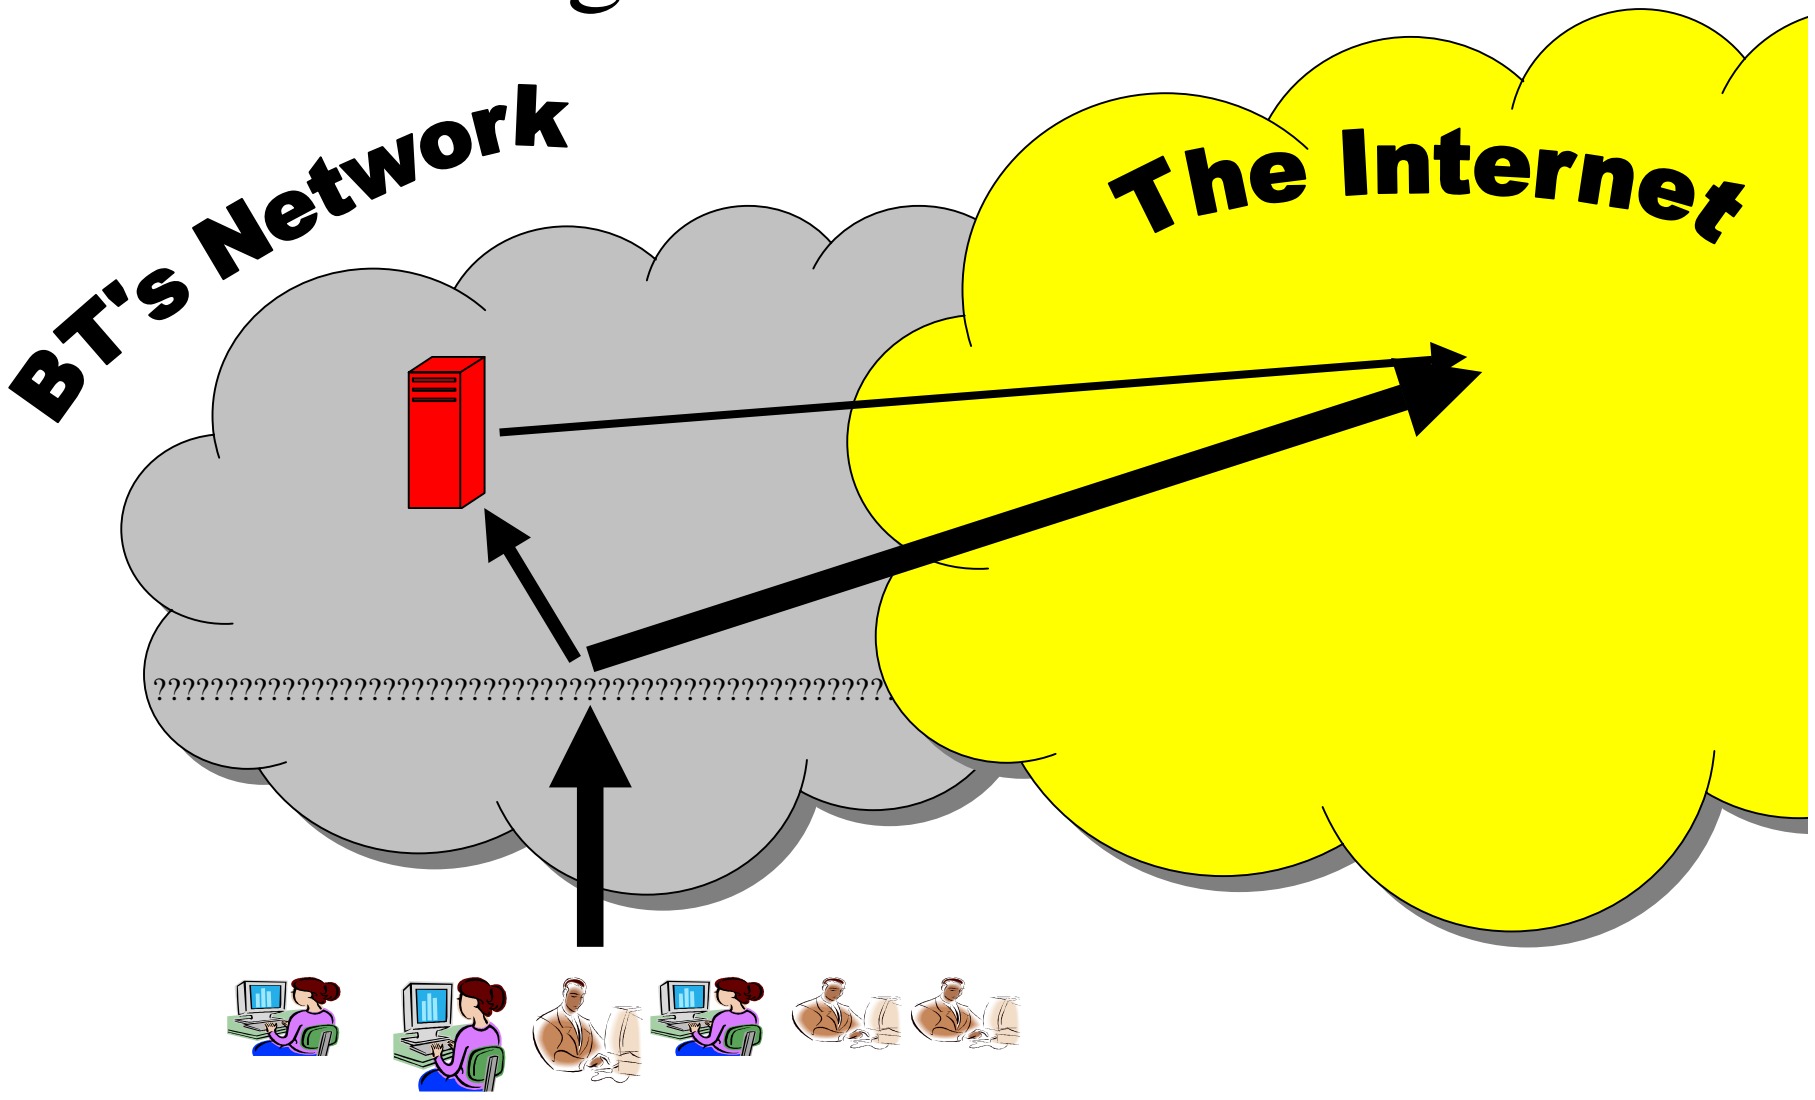

Fragility of Cleanfeed

- Evading either stage evades the system
  - all previous attacks continue to be relevant
- PLUS can attack the system in new ways
  - the credulous will fail to notice Google (or iTunes) IP addresses in DNS results for wicked sites, flooding the second stage with legitimate traffic
  - the clueless will fail to spot local IP addresses in DNS results and construct routing loops

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

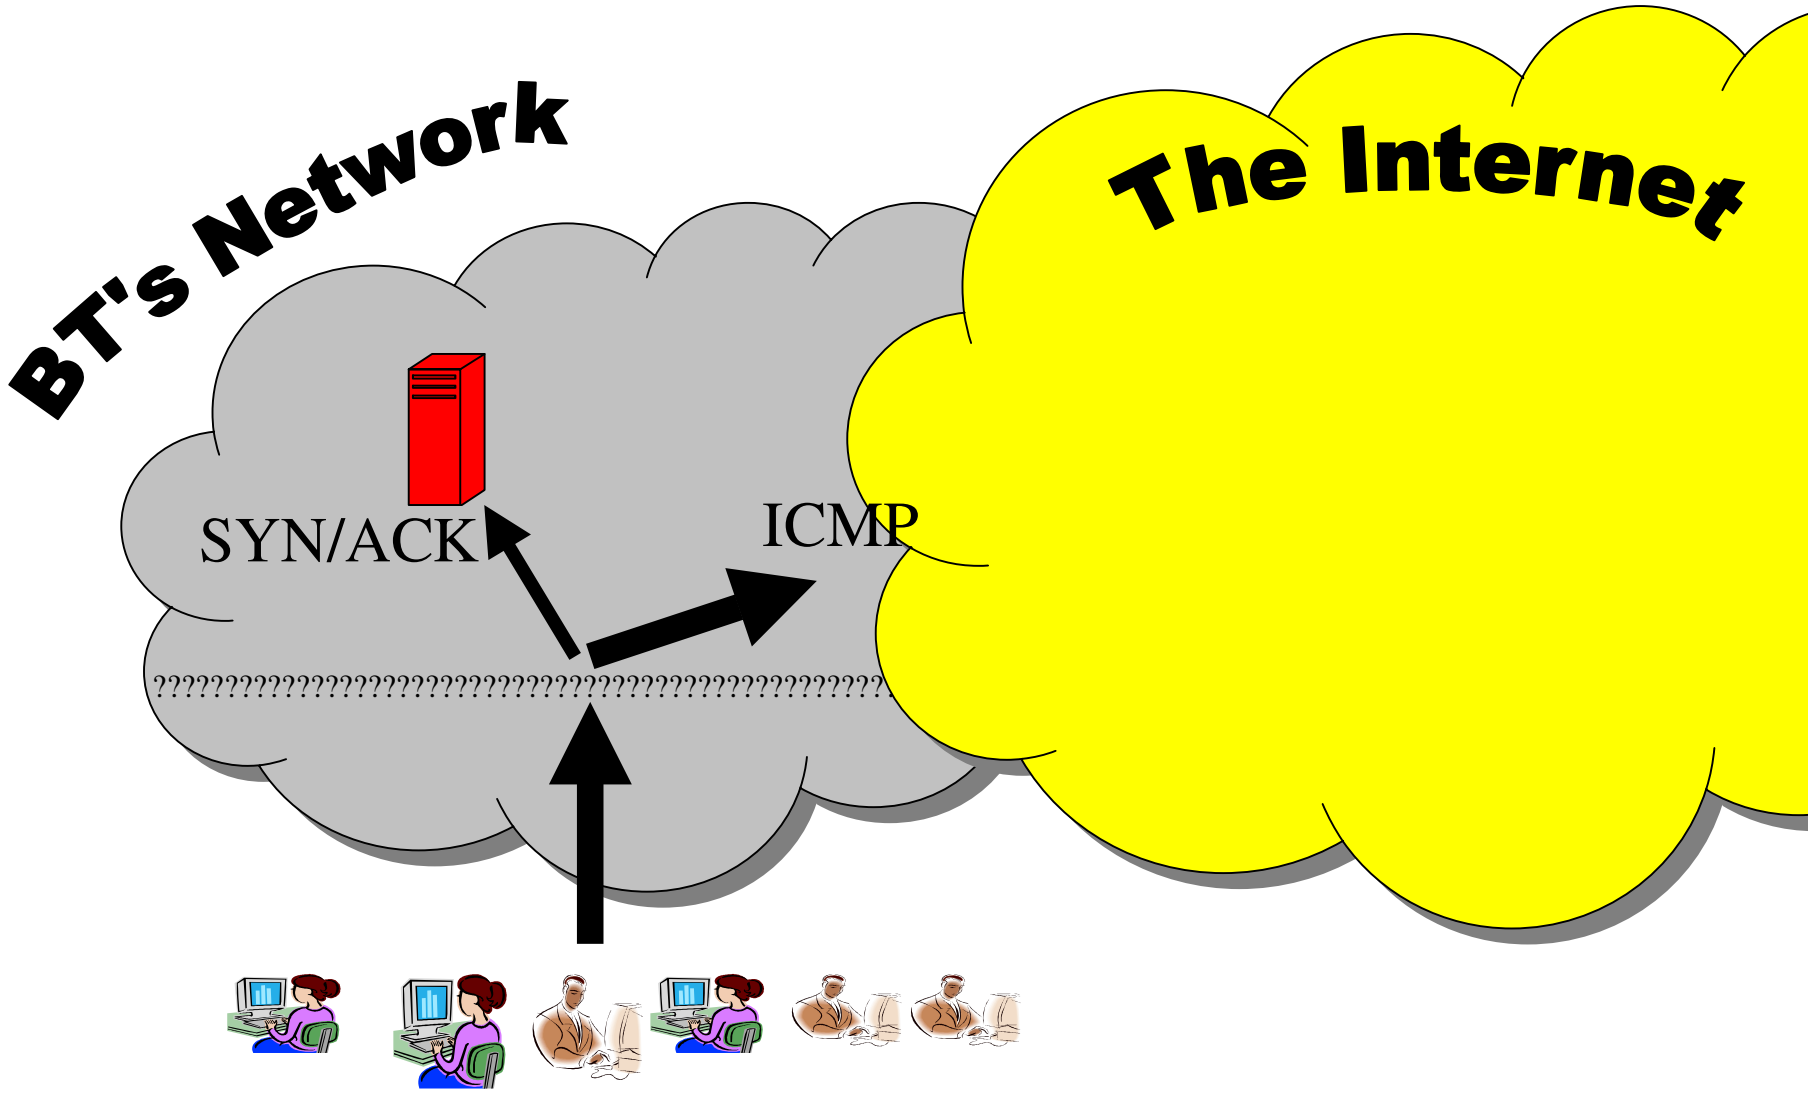

The oracle attack

- 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

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 many blocking systems

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 personal view, it is 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, allmymp3.com … etc etc etc

Conclusions

- Three basic ways of blocking content
- All have problems and can be evaded
- Hybrid systems can be lower cost, but have some extra problems as well
- Government signalling that blocking of sites on IWF list to become *de rigeur*
- Top of a very slippery slope

ISP Content Filtering: methods, failures and some politics

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

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

UNIVERSITY OF CAMBRIDGE
Computer Laboratory