“Proof-of-Work” Proves Not To Work

Richard Clayton

(joint work with Ben Laurie)

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Economics/Networks/Security
University of Cambridge
Seminar: 7th June 2004

Summary

- The current “spam” problem
- Viewing spam as an “economic” problem
- Proof-of-work mechanisms
- How much proof do you want?
- Analysis from an economic viewpoint
- Analysis from a security viewpoint
- Conclusions

The ages of “spam”

- Clueless sales & marketing personnel
- Disposable dial-up accounts
- Open SMTP relay “rape”
- Broadband and open proxies
- Spam friendly trojans (sent via virus?)
- Brute Force password guessing

… and doubtless more tomorrow

Worldwide Spam Category Data May 2004

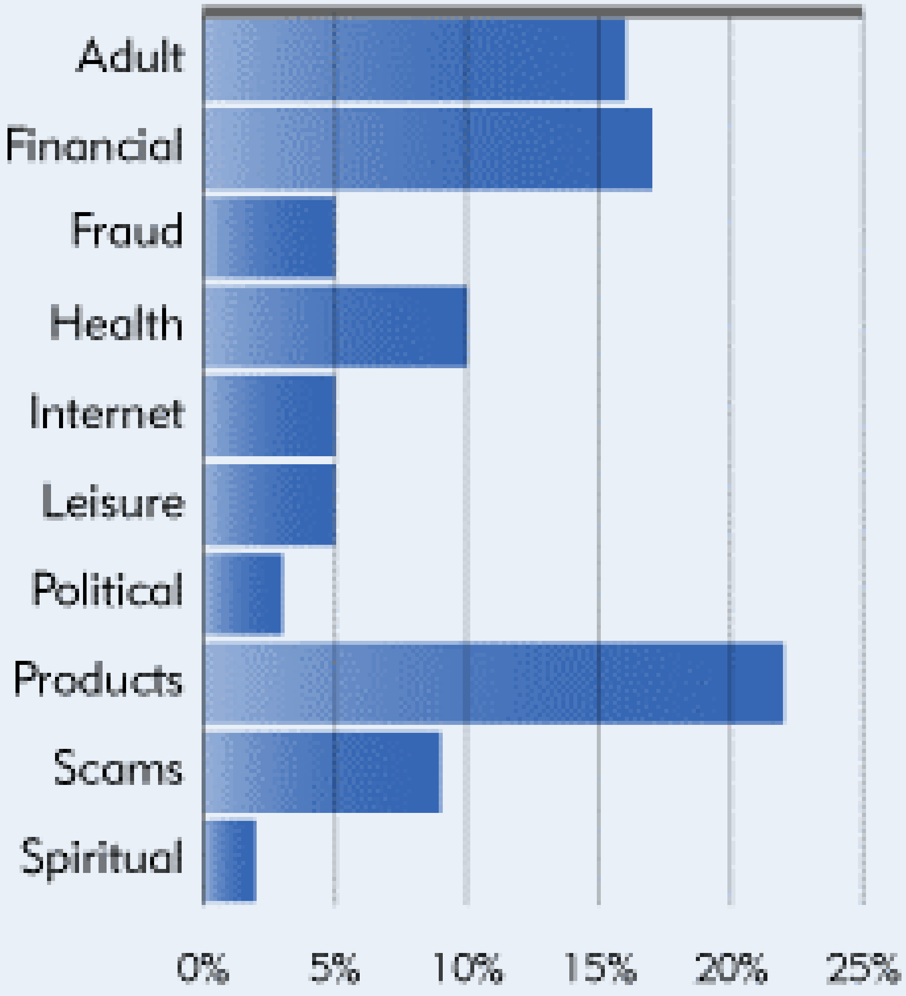

All other email attacks = 6%

Source: Brightmail Logistics and Operations Center

Percentages of Total Internet Email Identified as Spam

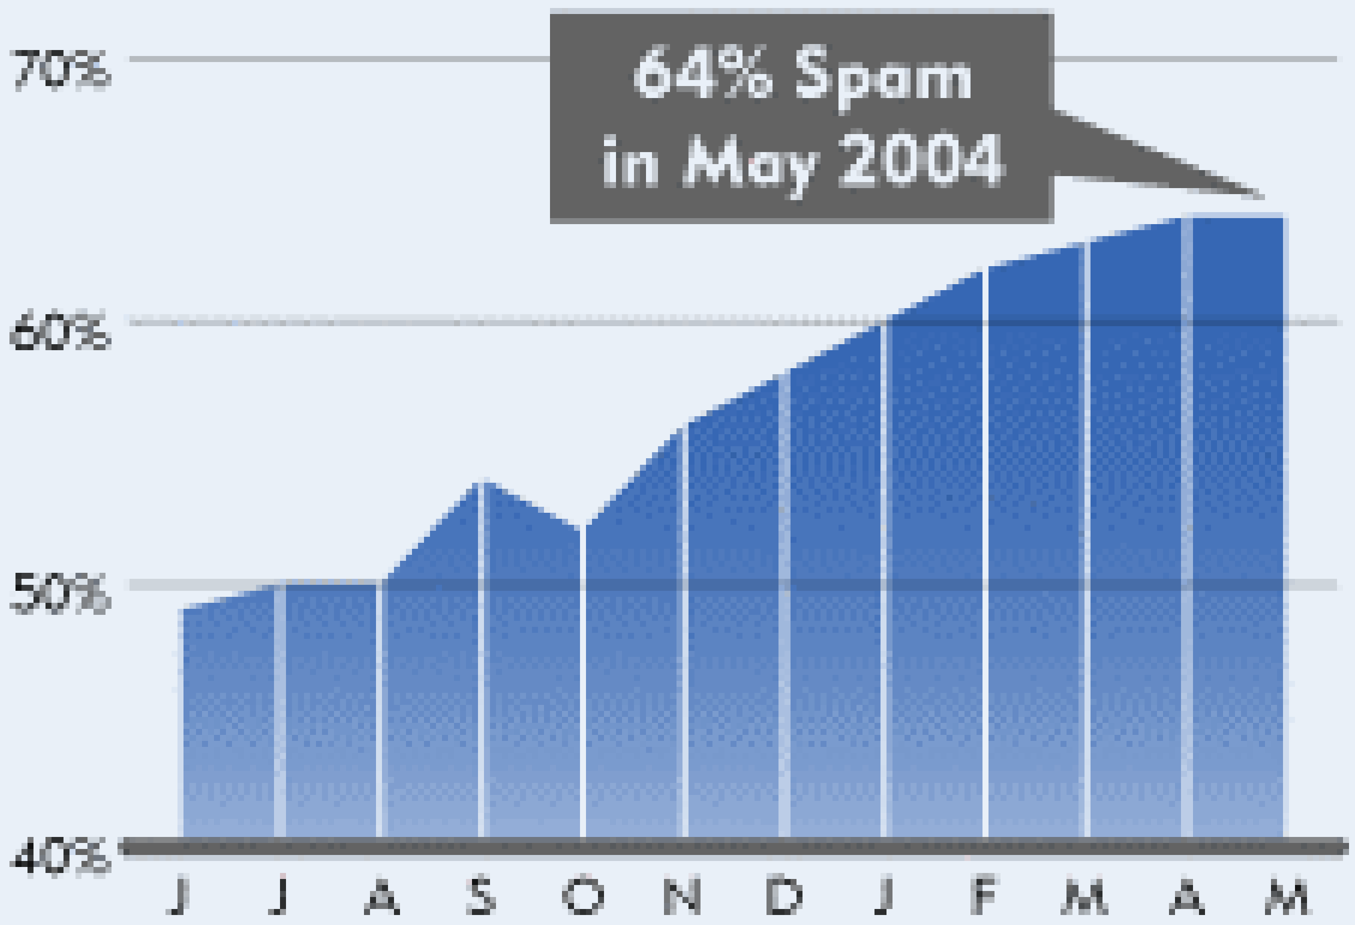

64% Spam
in May 2004

Over 100 Billion Email Messages
Filtered by Brightmail in May 2004

You think you get a lot of spam?

- junkname @ highwayman.com
  - May 2004: ~60,000 per day
- richard @ various domains (demon, turnpike etc etc)
  - 270 per day
- richard @ locomotive.com (last used Summer 1994)
  - 390 per day

Why?

- More multiply addressed spam
  - seems to be a policy change by the senders
  - this affects my counts, but not overall traffic
- More senders
  - SpamHaus lists 200+ major league spammers
- I’m an early adopter
  - my name will be on more lists
  - and lists come mainly from other lists

Let’s build “Something Else”

- Why should email be push not pull ?
  - actually on POP3 it’s pull already
- Doesn’t really tackle the human attention issue (how do you decide what to pull?)
  - It is not the Internet bandwidth cost that makes spam expensive!
- Main problem is that there’s very limited incentive to change to a new system

Countermeasures: Blocklists

- Idea is to record where spam comes from and then refuse to accept any more email from that particular source
- Usual implementation is using DNS queries
- Has scaled pretty well from initial ideas of a few dozen rogue sites
  - SORBS 1,414,266 open SOCKS proxies
    1,154,224 open HTTP proxies

Problems with Blocklists

- Many lists: no standard rules or processes
- Operators are pretty much unaccountable
  - SPEWS only reachable via *nan-ae*
- Have been used for personal vendettas
- Listing mail relays can be disproportionate
- Common to list /24s, affecting server farms
- Legacy lists (& shut-downs) are a problem

Countermeasures: Authentication

- Idea is to only allow authenticated senders to send you email
- Popular idea with Verisign, Microsoft and others who might handle the certi£icate$
- Essentially a cryptographically supplied whitelist (with a third party attesting to stranger’s probity)

Problems with Authentication

- Why should companies pay to send solicited email to their own customers?
- What happens when companies slip up?
  - how is the certificate be revoked?
- Spammers regularly compromise end-user systems – so will be authenticated anyway
- We’ve been authenticating IP addresses for years & it hasn’t been a silver bullet

Countermeasures: Filtering

- Idea is to assess content of email and decide that it is spam and discard it
- Works well for viruses
- Modern systems should not (!) suffer from the Scunthorpe effect
- Systems like SpamAssassin use a great many rules
- Currently this is fairly effective

Problems with Filtering

- False positives can cost the recipient dearly
- Legitimate email often blocked
  - eg opt-in promotional material
  - eg newsletters
  - eg airline ticket confirmations
- Spammers can use the filters too and tune their material to get through it
  - ie: spam is “evolving”

Is spam an Economics problem?

- Many argue that problem is “Economics”
  - – no charge for sending email
  - – hence “one in a million” response is profitable
- Hence the fix is to charge for email ?
  - – real money? 1p/email => $160 billion annually
    - phone companies would love this -- would we ?
  - – eCash? doesn’t seem to have happened yet !

Proof-of-work schemes I

- Idea is to show that you care enough about your email to have expended effort in doing a (rather pointless) calculation first
  - there are ideas for useful calculations eg “Bread Pudding Protocols” (Jakobsson & Juels 1999) but generally just warms up the planet ☹
- Original idea: Dwork & Naur : Crypto 1992
  - used central server ☹☹☹

Proof-of-work schemes II

- Reinvented as HashCash (Adam Back, 1997)
  - compute HASH(destination, time, nonce) such that result has “n” leading zeros
  - 2^n hard for sender, but trivial check for receiver
- Dwork, Goldberg, Naor (Crypto 2003)
  - analyse a function limited by memory speed
  - small variation between systems (factor of 4)
  - so this is much better than using classic HASH

Email statistics

- November 2003 (consistent stats available)
  - 2.30 x 10^8 Internet hosts (ISC)
  - 5.13 x 10^8 Internet users (Radicati)
  - 5.70 x 10^{10} emails sent daily (Radicati)
  - 56% of all email is “spam” (Brightmail)
- Hence the average situation is
  - 60 spam (& 50 real) emails per person per day
  - 125 real emails per host per day

What about “mailing lists” ?

- Expect to delegate proof-of-work analysis
- Lists common, but no published figures
- Inspected logs at large UK ISP (200K users)
    - this was after a spam filtering stage
    - consider identical source but >10 destinations
    - approximately 40% are of this form
- ie: reduce total to 75 emails per host per day
    - “back of envelope”, but only magnitude matters

How much work must we prove?

- Legitimate hosts must be able to send 75 emails per day (best case situation)
- Must reduce spam from 3.2 x 10^{10} per day
- Must allow for factor of 4 in capabilities
- Must assume spammers work 24 hours per day, but legitimate hosts may be switched off when not being actively used

… *so all we need to do is to pick "n"*

Economic analysis I

- Spammers charge 0.001 to 0.030¢ per email
  - survey in Goodman & Rounthwaite, 2004
- PC costs $500 / three years 50¢ per day
  - and pay electricity bill! 25¢ per day
- Spammer invests $50K and buys 100 PCs:
  - Salary $30K/annum 100¢ per day
  - So break-even at 35,000 emails/day/PC if can charge 0.005¢ each (ie: total 3.5 million /day)

[Scott Richter does 21 million/day @ 0.020¢]

Economic analysis II

- But spammers used to charge 0.1¢ per email (which leads to a break even rate of 1750)
- Spam response rates badly documented
  - Ms Betterly (WSJ Nov 2002) : 0.0023%
  - 0.0126% Iraqi Cards (“four times normal”)
- <u>If</u> 0.003% and 0.1¢ then cost of ads is $33/sale. Only viable for some products
  - $50/mortage lead; $85/cellphone, $60/pills

Economic analysis III

- Iraqi cards article (NYT 9 July 03) goes on:
  - best days: \$5000 profit per million emails
    ie: half a cent per email in commission
  - printer ink: \$500 to \$1200 per million emails
    ie: 0.05¢ to 0.12¢ per email in commission
- BUT note that legitimate email response rates are expected to be 0.7 to 1.6%
- Obviously wise to own more of value chain

Economic conclusion

- Good guys
  - – 75 emails/host (best case)
- Bad guys
  - – 1750 emails/host (if price returns to 0.1¢)
  - – but this will exclude low margin products ☺
- BUT bad guys have “factor of 4” advantage
- So some headroom here, but not lots & lots

AT CURRENT RESPONSE RATES

Security analysis I

- Lots of ***0wned*** machines out there
  - SORBS: 1.2M HTTP, 1.4M SOCKS proxies
  - Recent viruses have hit million+ machines each
- Currently easy to spot ***0wned*** machines
  - they send a lot of email!
- But what if they computed “proof-of-work”
  - quietly giving results to sender systems
  - hard to spot and so likely to be long-lived

Security analysis II

- Nov 2003, 3.2 x 10^{10} spam emails
- Suppose one million machines hijacked for proof-of-work (spammers share them out!)
- So, they only need to do 32,000 each
  - consistent with ISP figures for abused hosts
- If want 99% of our mailboxes to be "real" then must restrict spam to 250/host per day
- & for just 0.1% to be spam, then 25 per day

Security conclusion

- Good guys
  - – 75 emails/host (best case)
- Bad guys
  - – 250 emails/host (if spam is just 1% of mailbox)
- No “factor of 4” advantage this time
  - – unless spammers can choose *0wned* machines
- So **very** limited headroom
  - – & impossible to reach “one in a thousand” level

Real hosts : daily rates

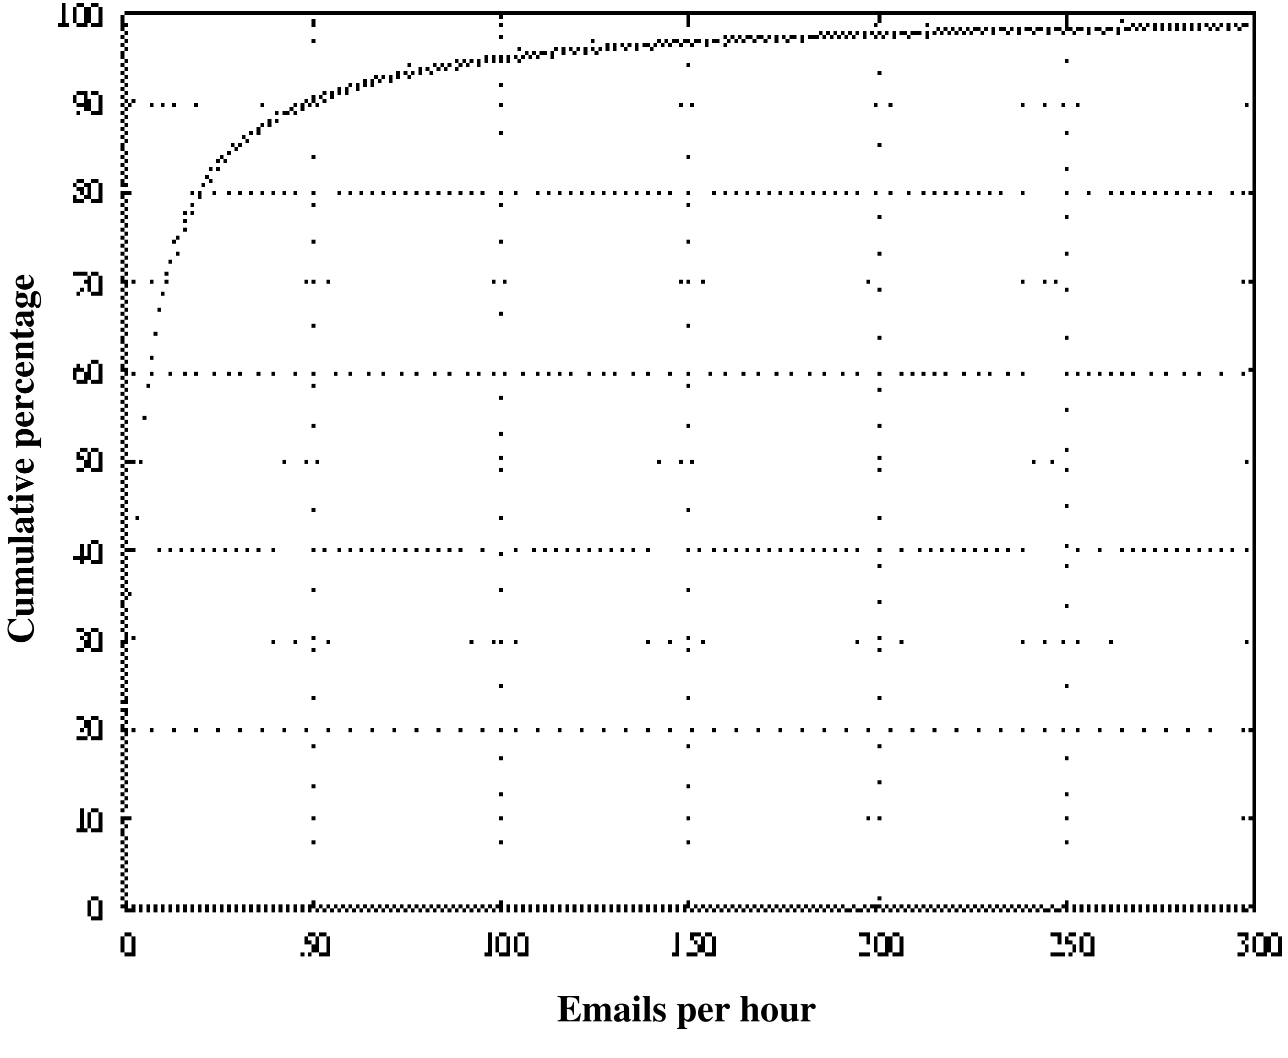

93.5% < 75

BUT

0.13% > 1750

1.56% > 250

viz: this impacts real senders

albeit some are just [exempted] mailing lists

Real hosts : hourly rates

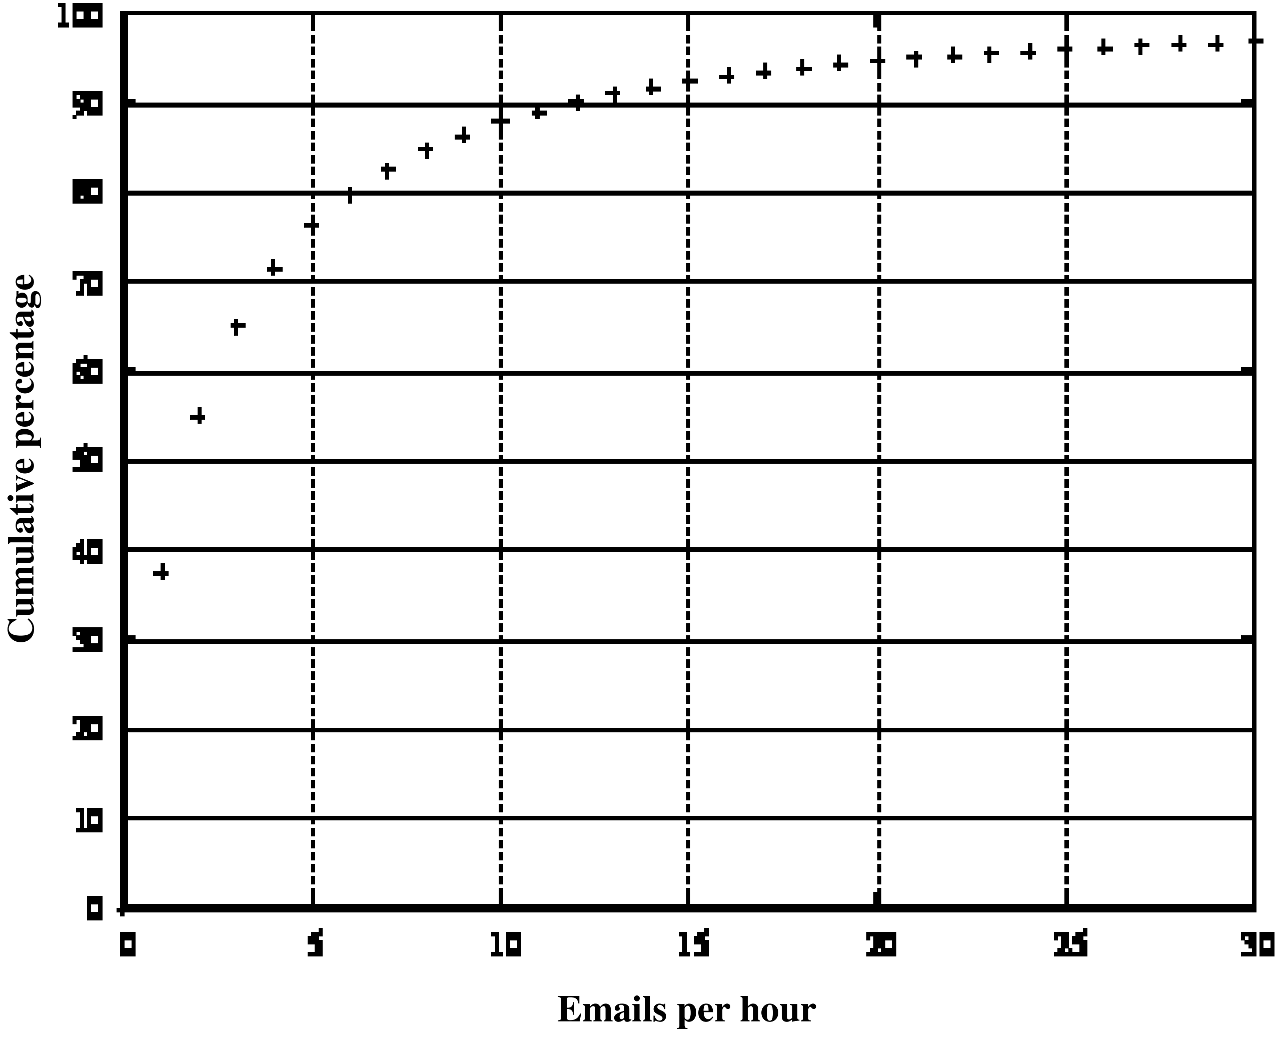

Spammers run
24 hours/day,
real users don't!

1% > 73/hour
i.e.1750/day

13% > 11/hour
i.e. 250/day

viz: this impacts
lots of people

Conclusions

- HashCash payment for email is attractive
- BUT spammer profit margins per sale mean that some will be able to afford the PCs to do the proof-of-work required
- BUT hijacking of end-user machines means impractical to restrict them to 1% of email
- Simplistic proof-of-work just doesn’t work!

“Proof-of-Work” ~~Proves~~ Not To Work
Proven

Ben Laurie — ALD — ben@algroup.co.uk

Richard Clayton — UNIVERSITY OF CAMBRIDGE Computer Laboratory — rnc1@cl.cam.ac.uk

& thanks to

Demon

The Cambridge-MIT Institute