

“Proof-of-Work” Proves Not To Work

Ben Laurie & Richard Clayton



**UNIVERSITY OF
CAMBRIDGE**
Computer Laboratory

Presented at: WEIS04
University of Minnesota
14th May 2004

Summary

- 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

“Spam”

- Unsolicited Bulk Email is a major problem
- Some 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)
 - 2.30×10^8 Internet hosts (ISC)
 - 5.13×10^8 Internet users (Radicati)
 - 5.70×10^{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 3.2×10^{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
 - 0.0023% reported (Wall Street Journal, 2002)
- If 0.0023% and 0.1¢ then cost of ads is \$4.35/sale. So viable for many products
 - \$50/mortgage lead; \$85/cellphone, \$60/pills
- Legitimate email response rates 0.7 to 1.6%

Economic conclusion

- Good guys
 - 75 emails/host (best case)
- Bad guys
 - 1750 emails/host (if price returns to 0.1¢)
- BUT bad guys will have “factor of 4” advantage over many good guys
- So some headroom, but not much

Security analysis I

- Lots of *Owned* machines out there
 - SORBS: 960K HTTP, 1.2M SOCKS proxies
 - Recent viruses have hit million+ machines each
- Currently easy to spot *Owned* 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

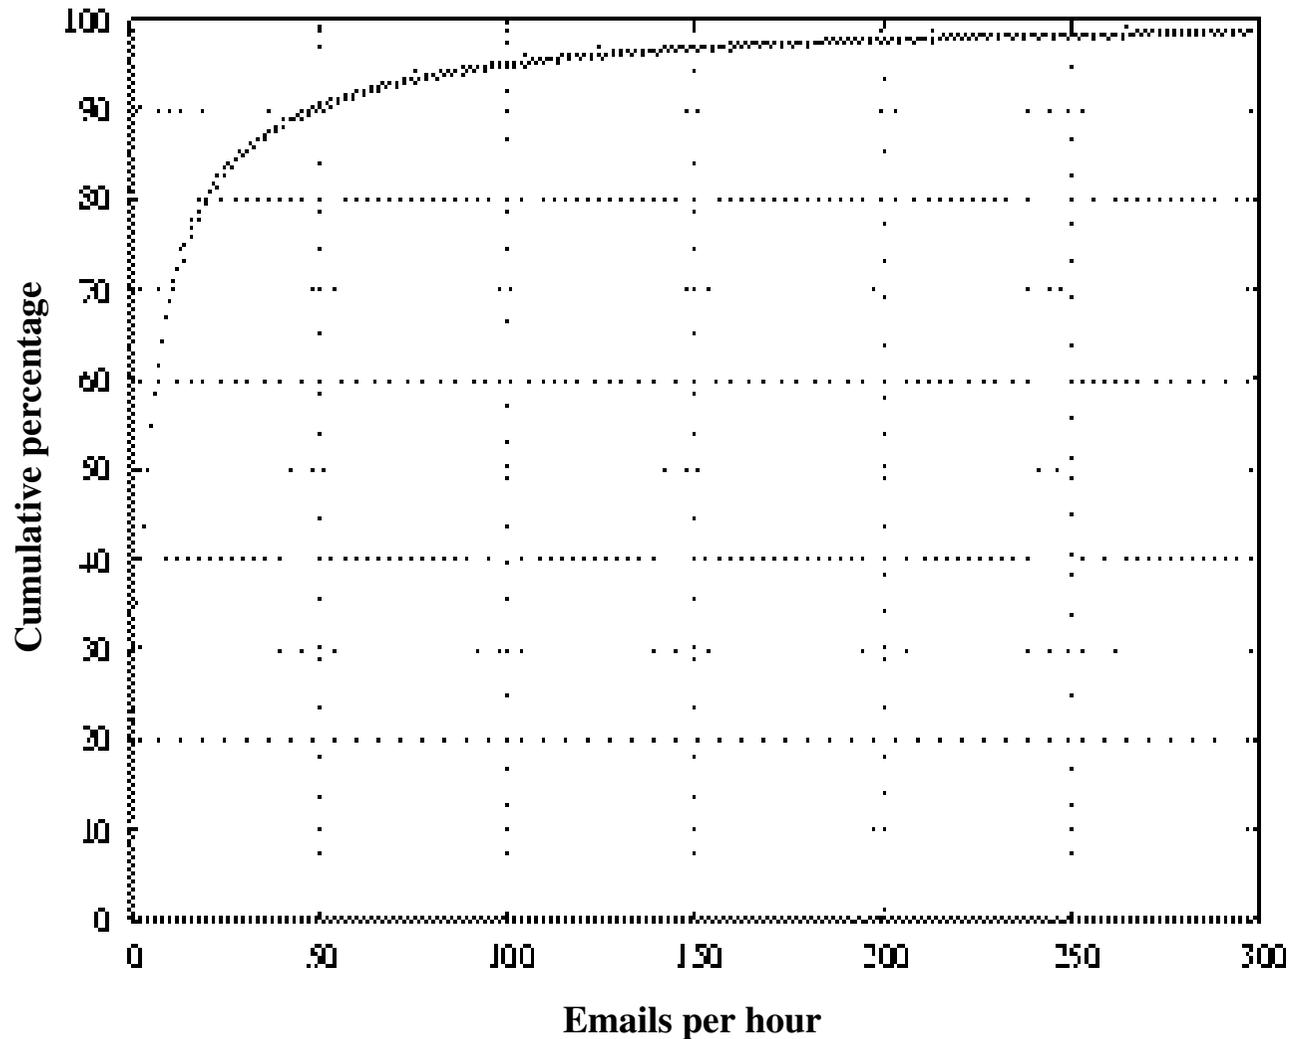
Security analysis II

- Nov 2003, 3.2×10^{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 *Owned* machines
- So **very** limited headroom
 - & impossible to reach “one in a thousand” level

Real hosts : daily rates



93.5% < 75

BUT

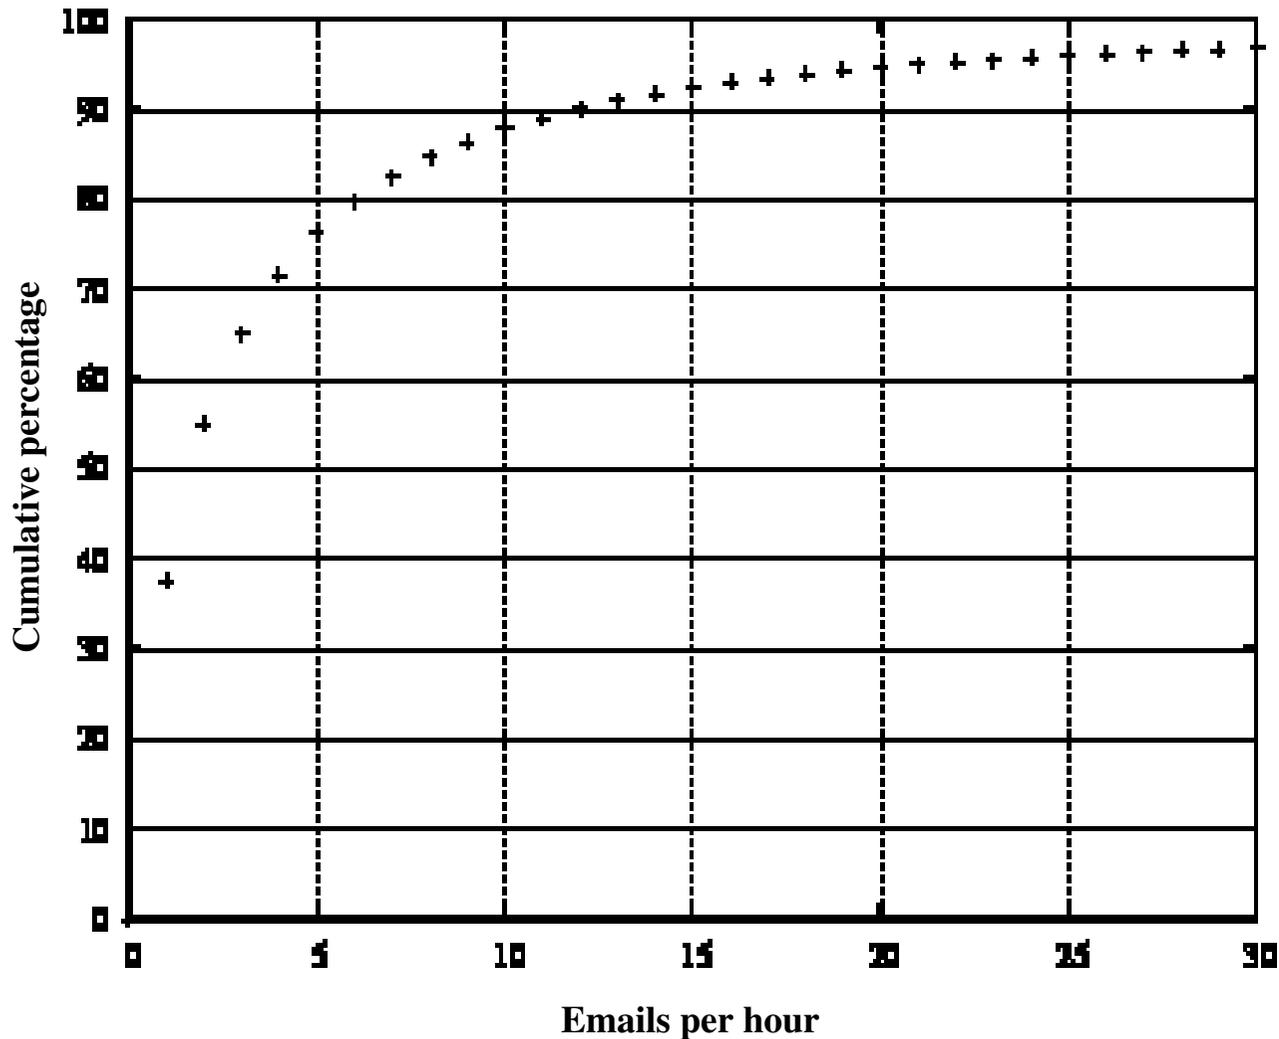
0.13% > 1750

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 they may well 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



ben@algroup.co.uk

Richard Clayton



UNIVERSITY OF
CAMBRIDGE
Computer Laboratory

rnc1@cl.cam.ac.uk

& thanks to

Demon



The
Cambridge-MIT
Institute