Modelling Incentives for Email Blocking Strategies

Andrei Serjantov

Richard Clayton

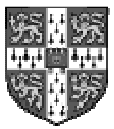

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

WEIS 2005
KSG, Harvard

2nd June 2005

Summary

- Setting the scene
- The model
- The implications of the model
- What is the pattern of outgoing email
- What is the pattern of incoming email
- Where next?

Setting the scene

- Email goes via ISP “smarthosts”
- Blacklists identify spam sources
  - may be a factor for Bayesian classifiers
  - may be used to block the sender altogether
- ISPs act in an ad hoc manner doing what seems to make sense to their sysadmins, and sometimes their customers
- Blacklists pretty much ad hoc as well!

The Model

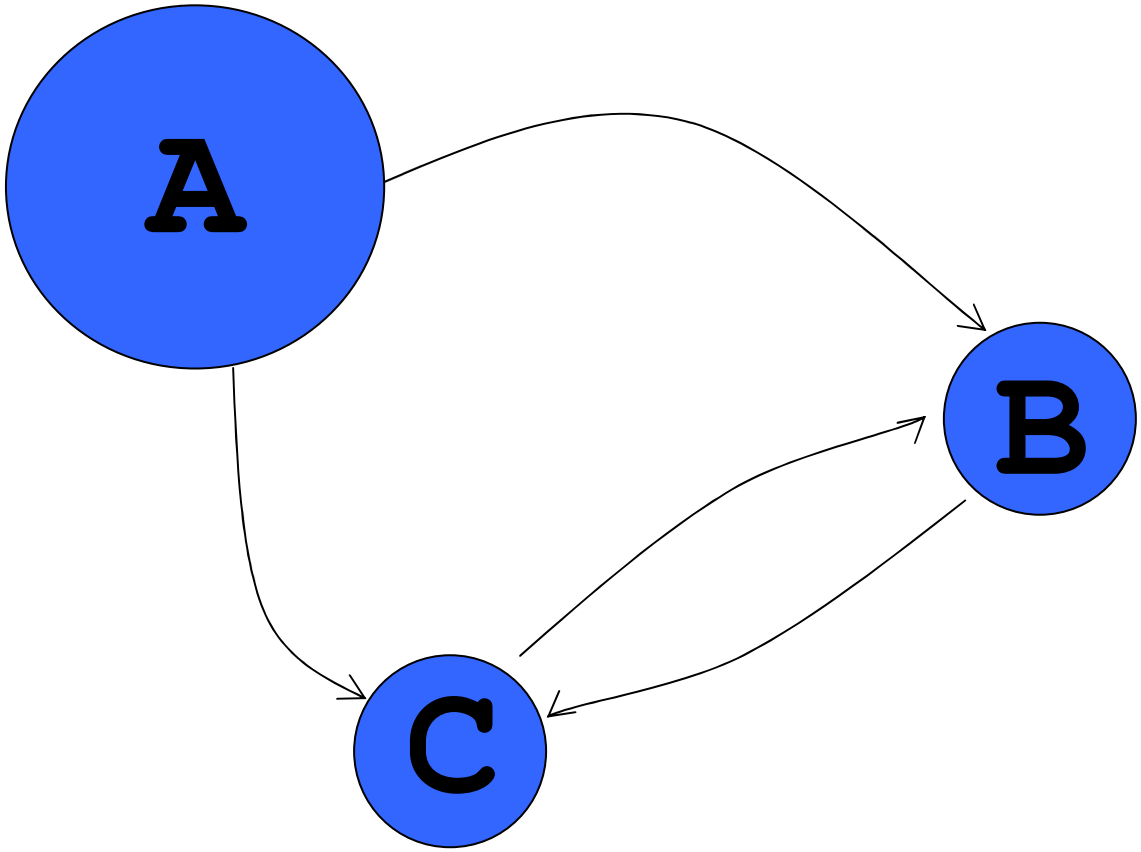

The Model

- Utility of ISP depends on its connectivity
  - Positive: ability to send email to others
    - Depends on how many people there are "out there"
  - Positive: reception of good email from others
    - Hard to perceive (all sorts of possible errors): ignore this term
  - Negative: reception of spam from others
    - Depends on how vulnerable remote clients are
    - And how many clients we have they may send to

$$Utility\,(A) = \left(\sum_{B} U(C_B)\right) - \left(\left|C_A\right| \times \sum_{B} V_B\left(C_B\right)\right)$$

Implications of the model

- The more “vulnerable” your clients are the bigger the negative term other ISPs see
  - they have to estimate this: guard your reputation!
- Dictionary attack spam affects large ISPs more (they have more clients who see it)
- Tit-for-tat blocking may work : remote ISP blocking us, we block them, our users don’t notice (!) but their users do

The view from large ISPs

- To large ISPs rest of world is very small
- Hence utility of connection to remote ISP dominated by how much spam they send
- Furthermore, utility equation dominated by self-sending term, and hence internal controls should be the overriding concern!

$$Utility\ (A)_{self} = U(C_A) - \left|C_A\right| \times V_A\left(C_A\right)$$

Outgoing email

- Measured outgoing email from Demon Internet (medium sized UK ISP) for four week period in March
  - excluded virus infected, spam sources etc
  - 82 000 customers (>50% use Hotmail etc)
  - 25 245 000 emails (of which 9 857 000 “bounces”)
  - 378 821 destination MX servers
  - but 240 850 only used once (typos + spam rejects)

ISP	Rôle	emails	customers
messagelabs.com	Spam filtering	1,361,916	35,641
hotmail.com	Global webmail	1,320,900	43,350
aol.com	Global ISP	820,645	37,674
btinternet.com	UK ISP	809,367	39,048
yahoo.co.uk	UK portal	367,327	24,302
demon.net	UK ISP (self)	363,112	15,212
ntl.com	UK ISP	337,441	25,174
yahoo.com	Global portal	298,491	18,139
uk.tiscali.com	UK ISP	235,858	22,022
virgin.net	UK ISP	189,389	18,358
schlund+partner	German web hosting	166,077	14,540
nhs.uk	UK health service	160,793	9,816
blueyonder.co.uk	UK ISP	149,521	14,677
pipex.net	UK ISP	97,495	9,576
spicerhaart.co.uk	UK estate agent	85,425	144
clara.net	UK ISP	82,309	8,106
mailcontrol.com	Spam filtering	80,941	6,978
global.net.uk	UK ISP	77,957	10,586
plus.net	UK ISP	77,080	9,289
postini.com	Spam filtering	74,777	6,092

Destinations: amount of email

- Power law distribution
  - see paper for straight line graph
- viz: same amount of email being sent to top 10 sites as to the next 100 as to the next 1000 as to the next 10000…
- A strategy that keeps only 10 destinations sweet (or only 100 etc) will fail

Destinations : number of senders

13 sites >10,000 customers sending to them

213 sites >1,000 customers sending to them

2601 sites >100 customers sending to them

- Potential for many complaints if just one of many other ISPs blocks Demon's email
- How much should Demon spend on their abuse team ?
  - clearly has a simple answer: Enough!

Incoming email

- 14 days incoming email
- 55.6 million emails
- 66.5% categorised as spam by “Brightmail”
- 13,378 sending ASs
- If an AS sent nothing but spam then would be rational to bar them
  - early test: one AS sent 9948, all spam in a day

Incoming: results inconclusive

- Many sources sent mainly spam, but still a few a day that were not
- Large volumes of spam (which would make real difference) accompanied by large volumes of good email
- Much more study needed
  - results much influenced by Brightmail
  - fast responses needed (infamous AS now OKish)

Conclusions

- Model explains much real world behaviour
- Figures clearly show very diverse aspect to communications: so ISPs cannot operate on a handful of special relationships
- Barring incoming email without impacting real traffic doesn’t look simple
- Still believe rational strategies are possible

Modelling Incentives for Email Blocking Strategies

Andrei Serjantov
Richard Clayton

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

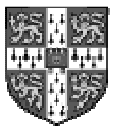

UNIVERSITY OF CAMBRIDGE
Computer Laboratory