

Spam

Edward W. Felten
Dept. of Computer Science
Princeton University

Scope of the Problem


- About 60% of all email is spam
 - Much is fraudulent
 - Much is inappropriate for kids
- 5% of U.S. net users have bought something from a spammer
 - Billions of dollars of sales
 - Spamming pays
- Will talk about email; but affects other communication technologies also

An Email Message

```
From: felten@cs.princeton.edu
To: lazowska@cs.washington.edu
Subject: mail forgery
Date: November 18, 2004
```

Actually, anybody can make a message like this. There's no inherent authentication of the receiver's address, and no guarantee that the message came from any particular place. Forgery is easy.

Email Transport



Complications: forwarding, mailing lists, autoresponders, etc.

What is Spam?

- (1) Email that the recipient doesn't want.

Problems:

- only defined after the fact
- ban raises First Amendment issues

- (2) Unsolicited email.

Problem: lots of unsolicited email is desired

What is Spam?

- (3) Unsolicited commercial email.

But what exactly does "unsolicited" mean?

Free Speech Issues

- Law sometimes allows speech, even when the listener doesn't want to hear it.
- Commercial speech less protected than political speech.
- At the very least, let's not block a message if both parties want it to get through.

Working Definition of Spam

- Any commercial, non-political email is spam, unless
- (a) the recipient has consented to receive it,
 - (b) the sender and receiver have an ongoing business relationship, or
 - (c) the message relates to an ongoing commercial transaction between the sender and receiver.

Note: just looking at a message won't tell you whether or not it's spam.

Anti-spam Measures

- Enforce laws against wire fraud, false medical claims, etc.
- Require accurate labeling of origin; allow filtering by origin
 - Big spammer just sentenced to nine years in VA state prison for mislabeling

Private Law suits by ISPs

- ISP sends spammer cease-and-desist letter
- Spammer keeps sending spam
- ISP files suit
 - Claiming cyber-trespass
 - Seeking money damages
 - Seeking injunction against further spamming
- Some success so far, but mostly useful as deterrent

Blacklists

- Make list of known email addresses, or known IP addresses, of spammers
- Discard email from those addresses
- Problems
 - Spammers try to mislead about message origin
 - Spammers move around a lot
 - Innocent users sometimes end up sharing addresses with spammers
 - False accusations

Whitelists

- Make list of people/places you want to get email from
- In practical to accept email only from these people
- But still useful
 - Make other anti-spam measures more stringent
 - Exception for people on whitelist

Payment

- Try to raise cost of sending email
 - Ideally, raise more for spammers than for normal senders
- Pay in the form of:
 - Money
 - Wasted computational resources
 - Human attention

Problems with payment

- If using real money, involves the banking system
- If paying in resources, waste of resources
 - Resources are cheap for spammers anyway
- Deters some legitimate email - especially big (legitimate) mailing lists

Sender authentication

- Various schemes
- Make sure that email comes from the right place, given the (claimed) sender
 - e.g. mymail.com from a Princeton IP address
- Works okay, but
 - Complicated in presence of forwarding etc.
 - Doesn't address spam bots on stolen machines

Content-Based Filtering

- Classify incoming messages based on contents
 - Apply fixed rules (e.g. Spam Assassin)
 - Machine learning, based on user labeling
 - Word-based Bayesian learning

Filtering Issues

- Fairly accurate, but not foolproof
 - Trade off false positives vs. false negatives
 - Still need to look at suspected-spam messages
- Spammers using countermeasures
 - "word salad"

Case Study: Do-Not-Email List

- In CAN-SPAM Act, Congress asked FTC to study a National Do-Not-Email (DNE) list
 - Like Do-Not-Call list for telemarketing
- Congress asked:
 - Should we have a DNE list?
 - If we have one, how should it work?
- FTC hired experts (including me) to give technical advice.

DNE List: Law

- Users can put their email addresses on the DNE list.
- Domain owner can put whole domain (e.g. washington.edu) on DNE list.
- Illegal to send spam to anybody on the list.

DNE List: Approaches

- Give spammers the list
 - Very bad idea: "whom-to-spam" list
 - Can seed each spammer's list with "telltale" addresses? (interesting CS theory problem.)
- Spammer submits their mailing list to DNE service; service returns "scrubbed" list
 - Spammer still learns about some valid addresses
 - Might be able to limit this by limiting access, charging for access, etc.

DNE List: Approaches

- Spam-forwarding service
 - Spammer must direct all spam through a DNE service
 - Service forwards email to addresses not on DNE list
 - Silently drops if address is on list
 - Doesn't leak information about list
 - Irony: as an anti-spam measure, the government is forwarding spam
- All approaches: risk that list will leak

Outlaw Spam

- Biggest problem for DNE List is outlaw spammers
 - Ignore the law
 - Send spam from stolen machines
 - Very hard to catch them

Spam: Bottom Line

- Spam will be with us, as long as people buy stuff from spammers.
- People will keep buying the kinds of products that spammers sell.
- At best, we'll fight to a stalemate.