Size Does Matter!

From the Age of Closed-Loop to the Age of Open-Loop

Roch Guerin
University of Pennsylvania
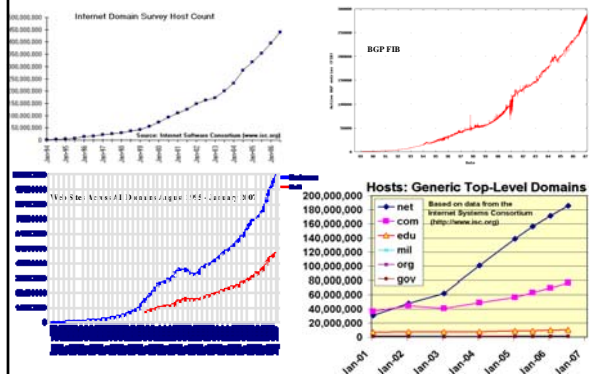

Premises

- The Internet has many flaws
  - Security
  - Predictability
  - Manageability
  - Etc.
- But it keeps growing at an unabated rate

2

And then some more…

So How Do We Control or Influence

- Something that already looked like this in 1999

- Grew to something like this by 2005, and keeps expanding!

4

Implications

- Whether we like it or not, the Internet has a "life" of its own and will continue expanding
  - It's hard to redesign or upgrade a fast moving train

- Scale makes most hard problems harder
  - Optimization, control, prediction, etc.
- But scale also offers many opportunities
  - Not everything goes wrong everywhere at the same time

5

Directions

- So what should we do if we want to make the Internet more robust and more reliable?
  - Don't try to control the uncontrollable
  - Devise solutions that exploit scale

- Solution space with two main components
  1. *Diversity* as a means to robustness
     - It's been pretty successful in other settings (physical layer)
  2. *Open-loop* approaches
     - Proactive rather than reactive

6

More Concretely…

- Having diversity is one thing, making it accessible and knowing how to use it is another
  - Realizing this has both control path and data path implications,

For instance:

- Control path
  - Routing should seek to maximize the number of available paths
    - Possibly distributing traffic unevenly across them as a function of "quality" (akin Cisco's EIGRP variance)
  - Some interesting challenges on both the algorithmic and scalability fronts

- Data path
  - Systematic redundancy and packet replication function in end-hosts and possibly routers
    - Packets from the same flow should not be sent on the same path
  - Decisions on how to leverage diversity should be mostly an open-loop process

7

Some Attempts at Justifications (1)
Generic Arguments

- A growing body of solutions that effectively leverage the Internet's diversity
  - Path switching
  - CDN and/or P2P overlays
- A resurgence of open-loop proposals
  - Oblivious routing
  - Diversity (multi-path) routing and coding
  - Multi-topology routing for standby backup paths
- With many approaches combining the two

8

Some Attempts at Justifications (2)
"Pet" Projects

- Improving throughput stability through path diversity and diversity coding
  - Exploring the trade-off of higher load versus increased success probability
- Algorithms for multipath maximization
  - Waste some bandwidth but provide as many backups as possible to everyone
  - Do it in a distributed way while preventing loops
- Leveraging multi-topology routing for joint performance and robustness optimization
  - Applicable to both intra and inter-domain routing

9

Implications for the Future Internet

- Hosts (or maybe access gateways) will
  - Encode all their data with some level of diversity coding (and obviously decode it…)
  - Distribute packet transmissions across multiple paths, when feasible, or ask the network to do it for them
  - Form groups of "trusted" peer relays to use as occasional backups
- Routers will
  - Compute multiple paths to each destination, including backups
  - Distribute packets across available paths
  - Be capable of "intelligent" packet replication

10

Implications for Future Internet Experimentation

- Experimental Internet platform should support
  - **[control path]** Deployment of routing protocol extensions (or new protocols) that enable computation and use of pre-computed backup routes
  - **[data path]** Large-scale exploration of the trade-off associated with packet replication (higher load vs higher resiliency)

Basically, what happens when "everybody" does it…

11

Sources

- Internet growth statistics taken from
  - http://www.potaroo.net/tools/asns
  - http://www.isc.org
  - http://navigators.com/stats.html
  - http://news.netcraft.com/archives/web_server_survey.html
  - http://www.opte.org/maps/

12