A Path Impairment Observatory

Brian Trammell and Mirja Kühlewind, NSG, ETH Zürich

HOPS(P)RG meeting, Yokohama, Monday 2 November 2015

The problem

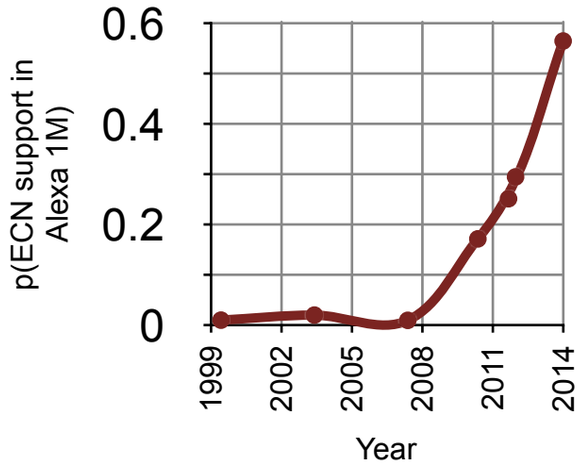

Modification	Planetlab	Ark
NAT	74.9%	79.0%
ECN IP	13.7%	13.2%
ISN	10.7%	1.8%
MSS	10.8%	5.9%
Exp. Option	8.8%	0.5%
MPCAPABLE	8.4%	0.3%
ECN TCP	0.6%	0.6%
SackOK	0.3%	0.0%
TS	0.3%	0.4%
WS	0.2%	0.2%

- What we want to know:
  - Given path *P*, does feature *f* work?
  - Given feature *f*, what is the probability it will work on a random path *P*?
- Lots of middlebox measurements and tools.
- But, the Internet is heterogeneous.
- So, no idea how comparable the results are.

[1]: B. Trammell et al. **Enabling Internet-Wide Deployment of ECN.** PAM, March 2015
[2]: R. Craven et al. **A Middlebox-Cooperative TCP for a non End-to-End Internet.** SIGCOMM, August 2014.

A solution

- Build an observatory which takes input from a diverse set of tools and measurement campaigns.
- Single observation {***t, P, c***} where
  - ***t****:* time at which the observation was taken (and assumed valid)
  - ***P****:* variable-precision designator of the path on which the observation was taken
  - ***c****:* variable-definition designator of the condition observed on that path
- Using path and condition characterization and equivalence operations, allows more formal comparison of diverse studies.

Requirements

- Accept data from a wide variety of sources, e.g.:
  - Raw output from tools we maintain.
  - Raw packet traces of active measurements.
  - “Here’s data from a measurement study, and references to commits in GitHub for the tools and configuration used to generate it.”
    - “Option foo breaks on these paths but not these paths; we’re not going to tell you where they are, but this set of them belong to a major mobile carrier.”
- Support path pseudonymization and aggregation for privacy.
- Support condition definition with enough precision to allow active measurements to reproduce observations on other paths.
- Integrate with existing tools, without restrictions on implementation.

Path designators

- Path designators ***P*** are variable-precision:
  - "The sequence of IP addresses $a_0,a_1,\ldots,a_{n-1},a_n$ between source a_0 and destination a_n"
  - "The source/destination IP address pair a_0,a_n"
  - "The BGP AS path $b_0,b_1,\ldots,b_{n-1},b_n$"
  - "The path between a_0 and a_n which contains a_1, b_3, and b_{n-1}"
  - "An access path with opaque identifier *i* with the following characteristics"
- Tradeoff between precision and measurement cost as well as privacy.
  - We expect much of the data accepted from third parties to be of the opaque identifier type.
- Set of operations for determining how equivalent two paths are (and therefore, whether one path answers to a query for another).

Condition expressions

- Ideal: a completely generic language for ***c*** describing patterns in common packet structures.
  - works for matching against arbitrary traces!
  - works for fuzzing/measurement packet generation and recognition!
  - works for automated tool test generation! but…
  - looks a lot like regular expressions.

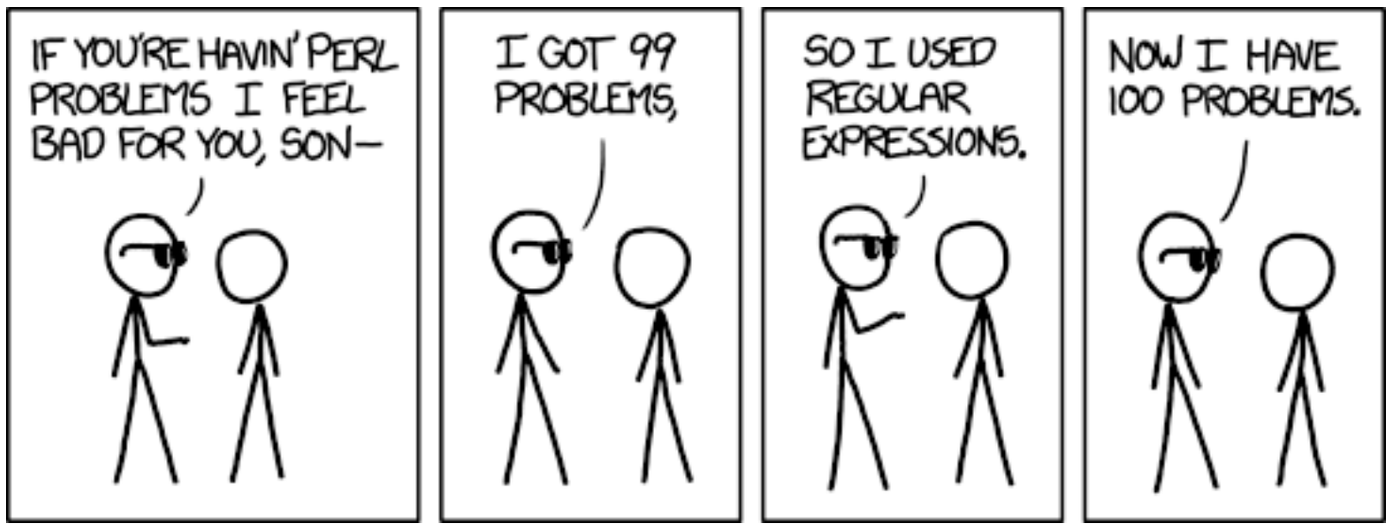

- Current work: TBD, support generic language(s) as well as code-reference.

image: Randall Munroe, XKCD 1171, cc-by-nc

Where do we go from here?

- We're just getting started:
  H2020 Measurement and Architecture for a Middleboxed Internet (MAMI): 2016-2018.
  - Initial implementation over a reasonable (O)RDBMS (PostgreSQL) by IETF 97 Seoul.
  - See https://github.com/mami-project for code.
- We intend to operate a long-term observatory during and after the scope of the project.
- Talk to us about how we can help each other.