Using Measurement Data to Construct a Network-Wide View

Jennifer Rexford
AT&T Labs—Research
Florham Park, NJ

http://www.research.att.com/~jrex

Executive Summary

- Key network operations tasks benefit from a *domain-wide* view of traffic and routing
- Router vendors should help us get it

For the folks in the back of the room…

… this portion is intentionally left blank

Motivating Applications

- Usage reporting/trending
  - Application mix (P2P, Web, DNS, etc.)
  - View per customer, per peer, etc.
- Detect, diagnose, & fix problems
  - Flash crowds, DDoS attacks, new hot apps
  - Route flaps, blackholes, highjacked prefixes
- Traffic engineering & capacity planning
  - Tuning routing configuration to the traffic
  - Adding new links, routers, peers, proxies, …
  - Predicting the effects of changes in advance

Measure, Model, and Control

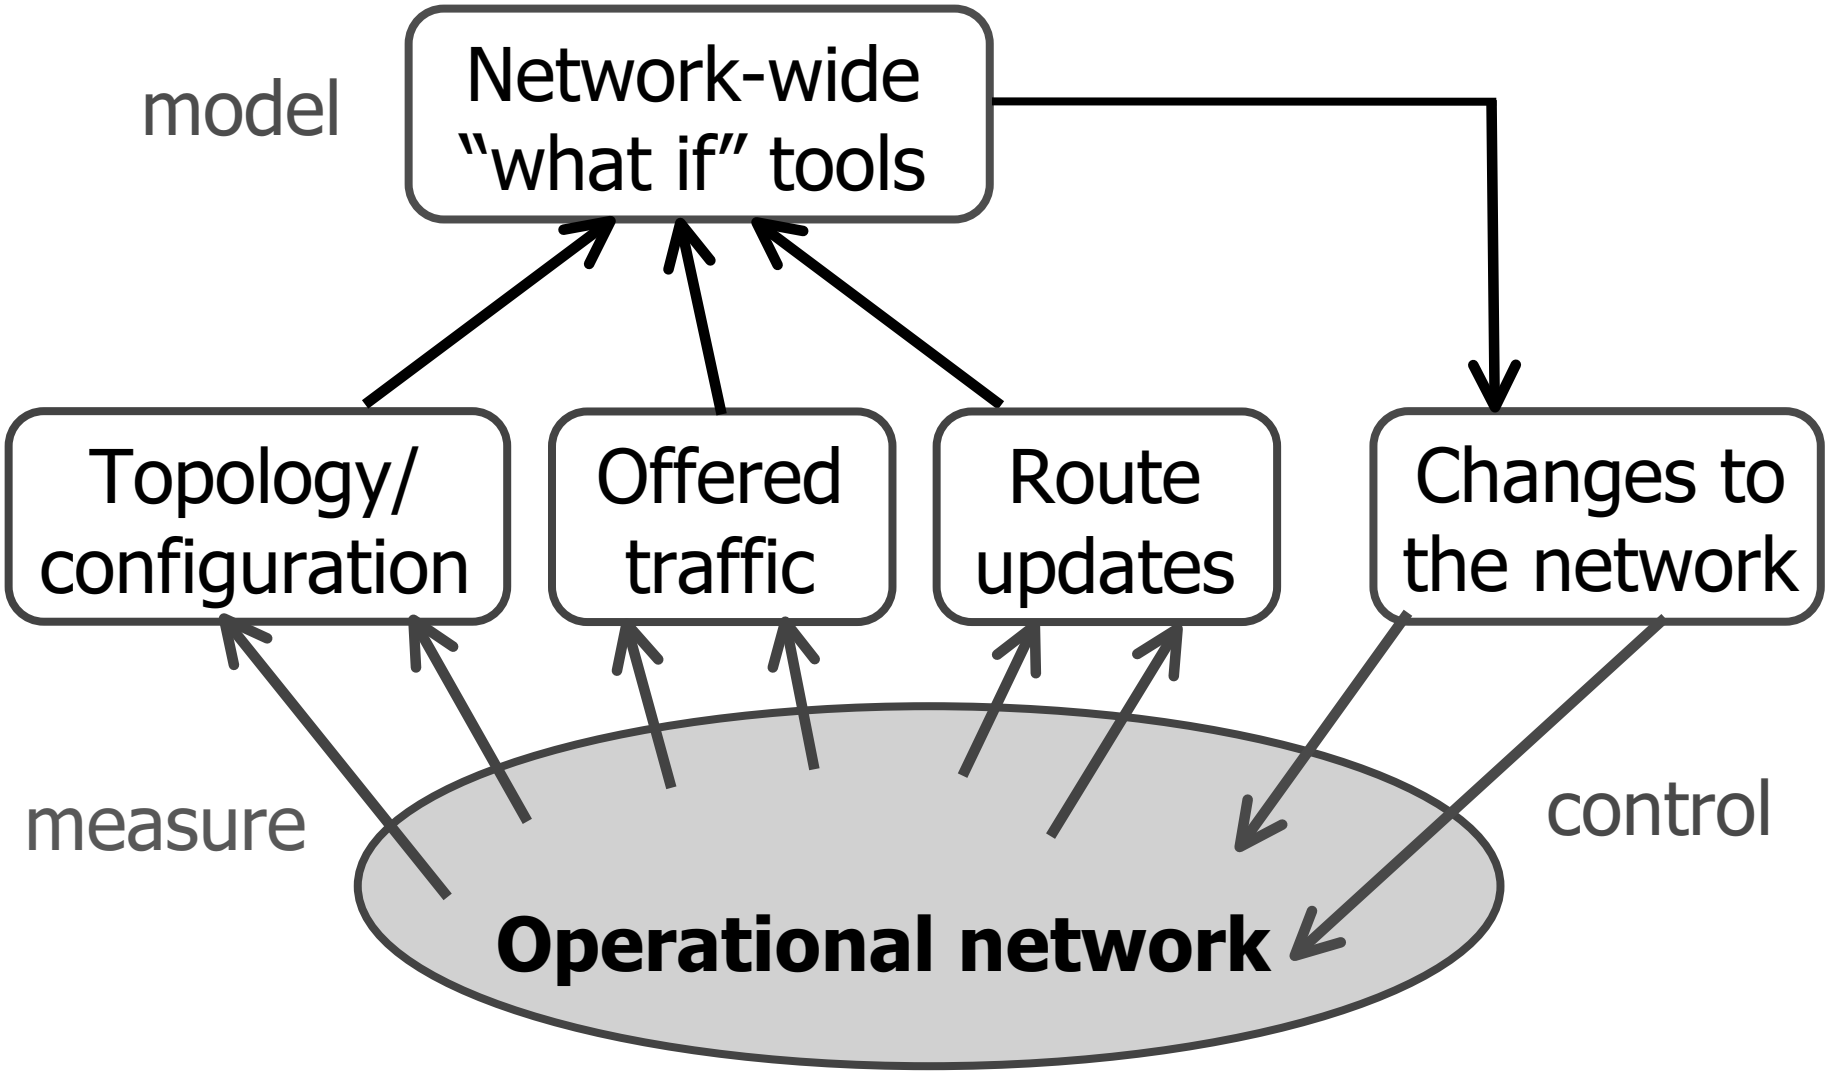

Network-Wide View

- Topology and configuration
  - Up routers, links, and routing sessions
  - Current IGP weights and BGP policies
- Traffic matrix/demands
  - Load between ingress and egress points
  - Traffic per destination prefix per ingress
- Routing advertisements
  - Routes learned via each eBGP session
  - All routes, before any import processing

Traffic: Packet Sampling

- IETF working group on packet sampling
  - Minimal measurement functionality
  - Suitable for high-speed line cards
  - Tunable overhead/accuracy trade-offs
- Basic idea: parallel filter/sample banks
  - Filter on header fields (src/dest, port #s, …)
  - 1-out-of-N sampling (random, periodic, hash)
  - Extract header fields, output link, IP prefix,…
  - Send group of records to a collection system

Psamp Functionality

On the line cards…

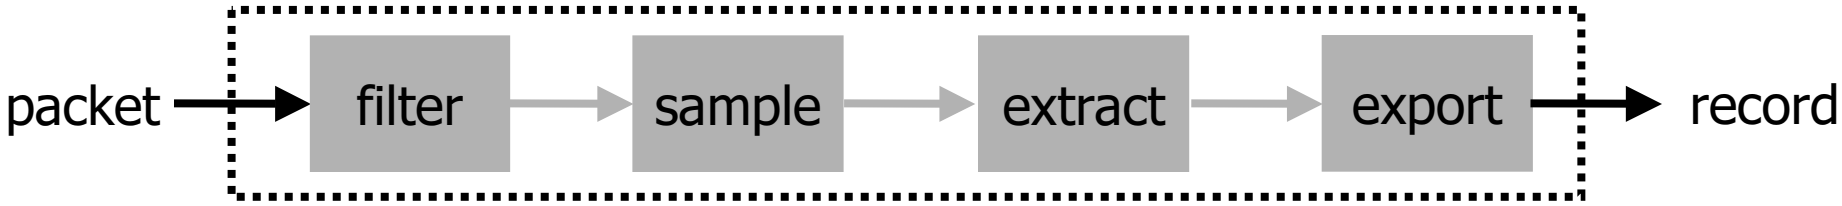

Example configurations

- Baseline: 1/10000 of all packets
- Customer: 1/100 on src/dest prefix
- DDoS: 1/100 on destination address
- Web: 1/1000 on port 80

I couldn't resist putting something here…

http://www.ietf.org/html.charters/psamp-charter.html

Routing: BGP and IGP Monitoring

- Periodic table dumps
  - Pro: all of the routes (best and alternate)
  - Con: coarse timescale, poor synchronization across routers, high router overhead
- BGP session with operational router(s)
  - Pro: continuous feed and limited overhead
  - Con: only best routes, only after import processing, resets of multi-hop session

Route Monitoring Session

- Onboard support for route monitoring
  - Special monitoring session with the router
  - Continuous export of received routing messages
  - Tolerance of transient reachability problems
- Data export
  - Concise format for higher efficiency
  - Omission of redundant info (e.g., refresh LSAs)

Conclusion

- Network-wide view for control
  - Topology, traffic, and routing
- Lightweight vendor support
  - Operation on high-speed links and routers
- Packet sampling (psamp)
  - Parallel filter/sample/extract/export banks
- Route monitoring session
  - Relaying of received routing updates/LSAs