Network Measurement

Jennifer Rexford

Fall 2010 (TTh 1:30-2:50 in COS 302)

COS 561: Advanced Computer Networks

http://www.cs.princeton.edu/courses/archive/fall10/cos561/

Measurement for Network Operators

Network Operations: Detecting the Problem

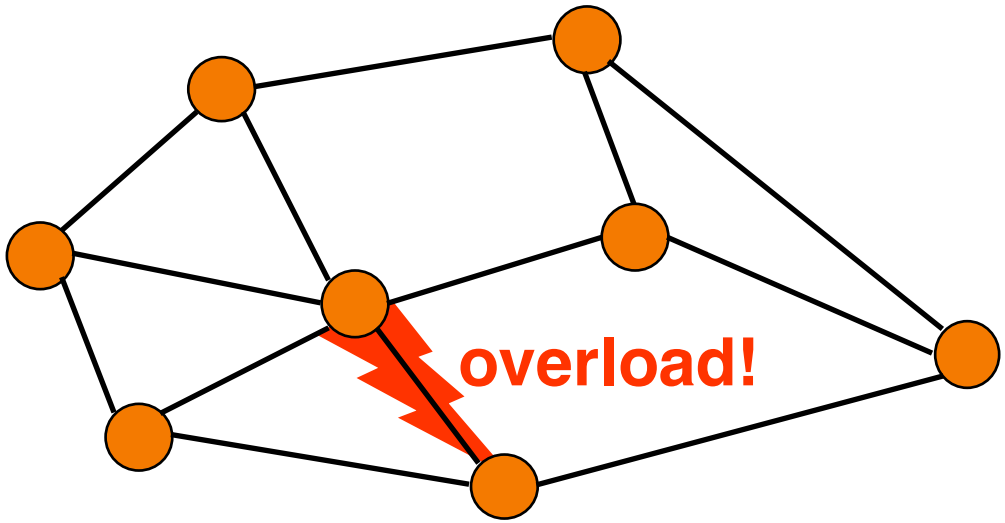

Detecting the problem!

- High utilization or loss statistics for the link?
- High delay or low throughput for probe traffic?
- Complaint from angry customer (via phone network)?

Network Operations: Excess Traffic

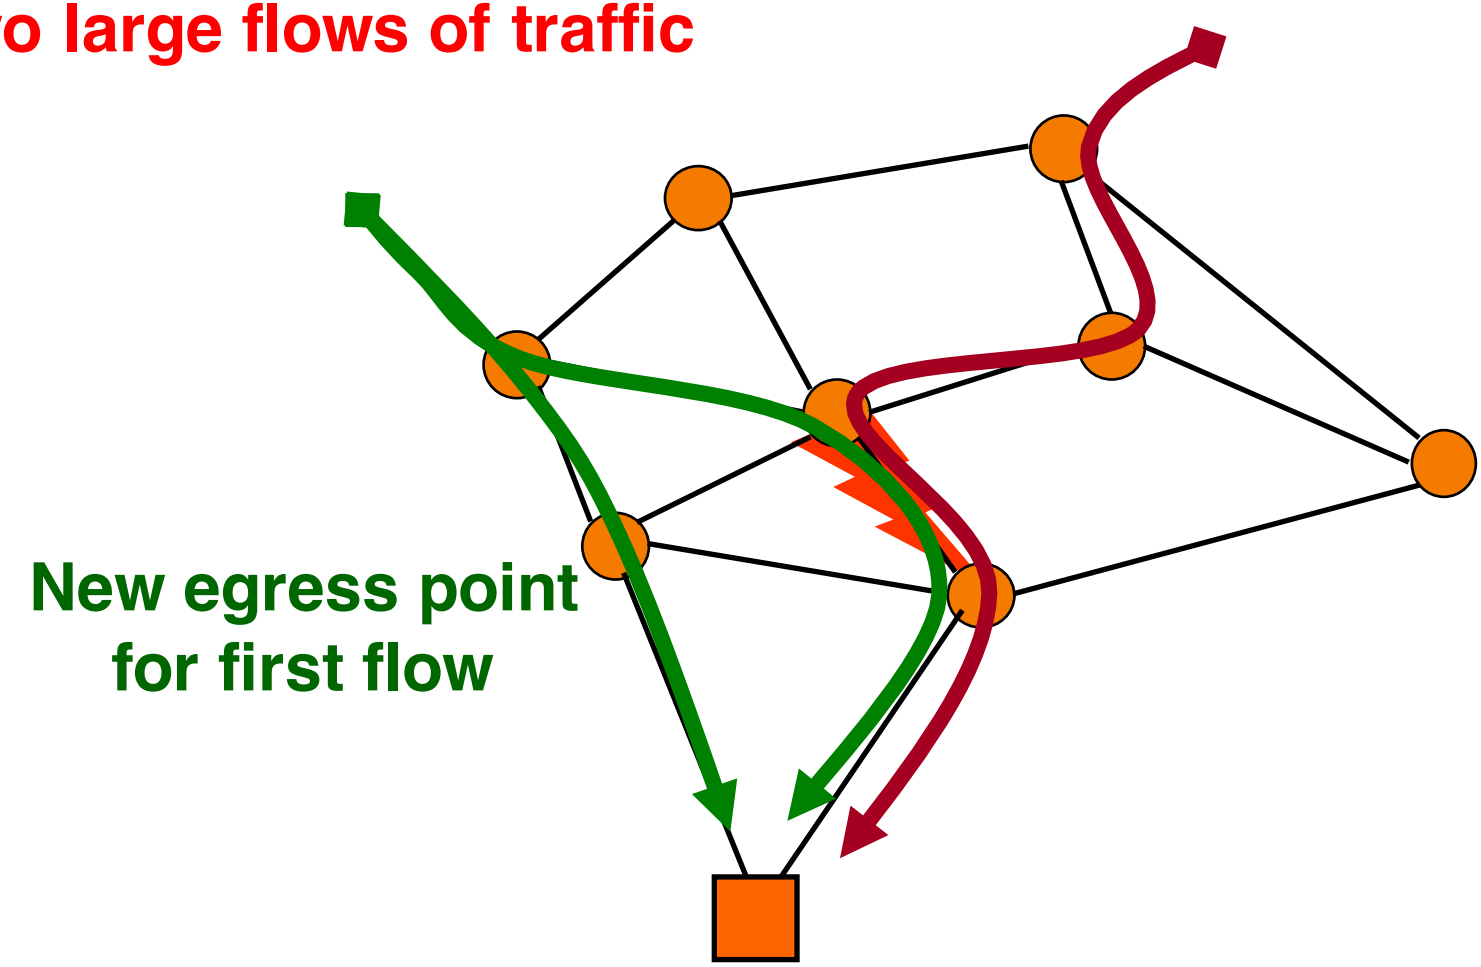

Network Operations: DoS Attack

Install packet filter

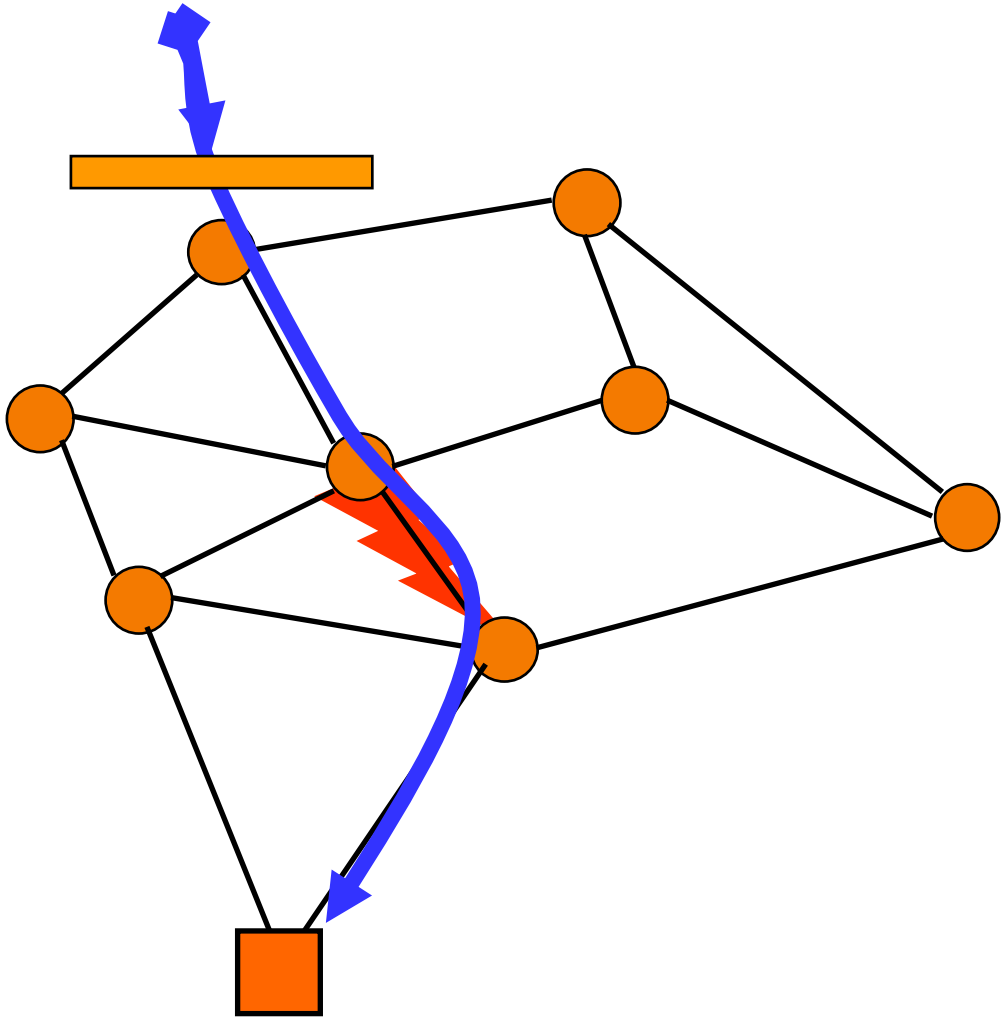

Web server back to life...

Network Operations: Link Failure

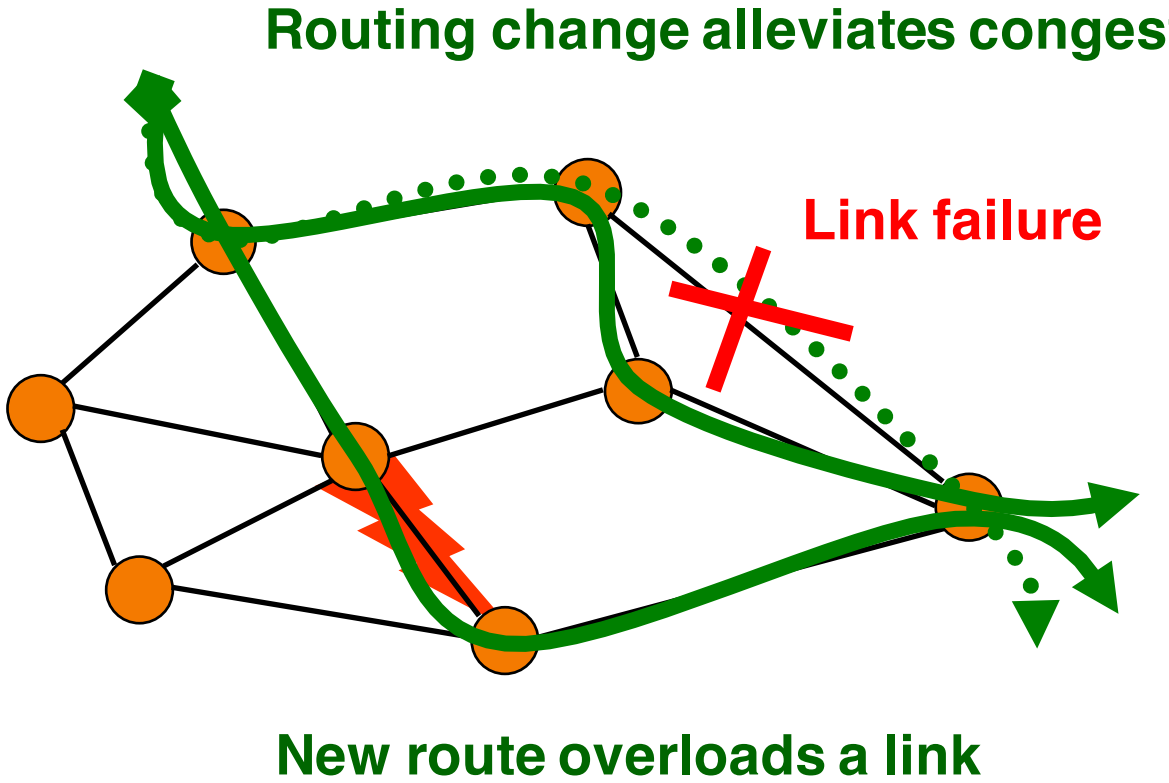

Summary of the Examples

- How to *detect* that a link is congested?
  - Periodic polling of link statistics
  - Active probes measuring performance
  - Customer complaints
- How to *diagnose* the reason for the congestion?
  - Change in user behavior
  - Denial of service attack
  - Router/link failure or policy change
- How to *fix* the problem, and model the effects?
  - Interdomain routing change
  - Installation of packet filters
  - Intradomain routing change

The Role of Traffic Measurement

- Operations (control)
  - Generate reports for customers and internal groups
  - Diagnose performance and reliability problems
  - Tune the configuration of the network to the traffic
  - Plan outlay of new equipment (routers, proxies, links)
- Science (discovery)
  - End-to-end characteristics of delay, throughput, and loss
  - Verification of models of TCP congestion control
  - Workload models capturing the behavior of Web users
  - Understanding self-similarity/multi-fractal traffic

Measurement for Network Operators

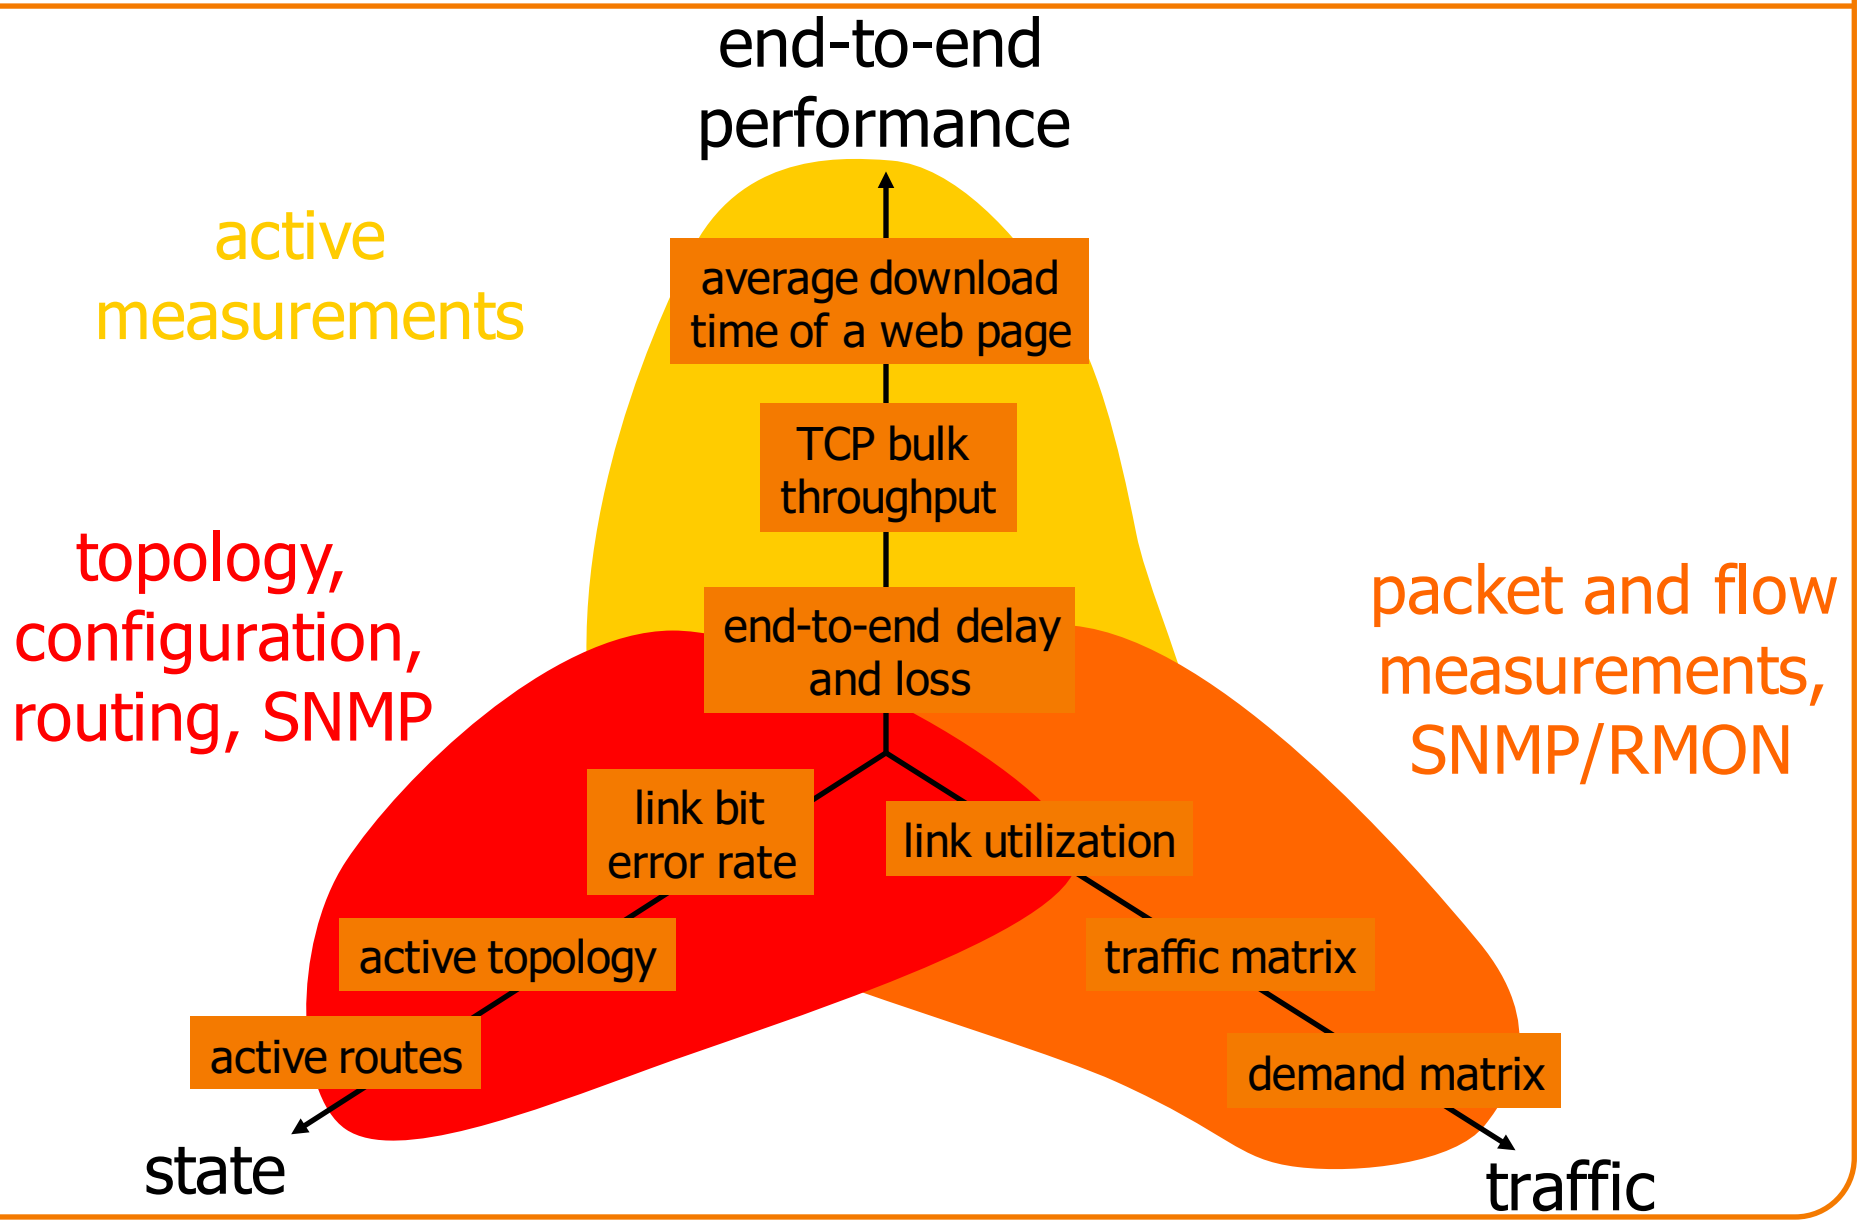

Measurement Challenges for Operators

- Network-wide view
  - Crucial for evaluating control actions
  - Multiple kinds of data from multiple locations
- Large scale
  - Large number of high-speed links and routers
  - Large volume of measurement data
  - High dimensionality of the data
- The “do no harm” principle
  - Don’t degrade router performance
  - Don’t require disabling key router features
  - Don’t overload the network with measurement data

Data Reduction Techniques

Streaming Algorithms that Perform Filtering, Aggregation, and Sampling

Controlling Measurement Overhead

- Measurement overhead
  - In some areas, could measure everything
  - Information processing not the bottleneck
  - Examples: geology, stock market,...
  - Networking: thinning is crucial!
- Basic methods to reduce measurement traffic:
  - Filtering
  - Aggregation
  - Sampling
  - ...and combinations thereof

Filtering

- Only record data for a subset of the traffic
  - Based on fields in the packet header
  - Ignore the rest of the traffic
- Examples
  - Matching a destination prefix (to a certain customer)
  - Of a certain service class (e.g., expedited forwarding)
  - Violating an ACL (access control list)
  - TCP SYN or RST (attacks, abandoned http download)

Aggregation

- Example: by communicating hosts
  - Independent variable: source-destination
  - Metric of interest: total # pkts, total # bytes
  - Variables aggregated over: everything else

src	dest	# pkts	# bytes
a.b.c.d	m.n.o.p	374	85498
e.f.g.h	q.r.s.t	7	280
i.j.k.l	u.v.w.x	48	3465
....	....	....	

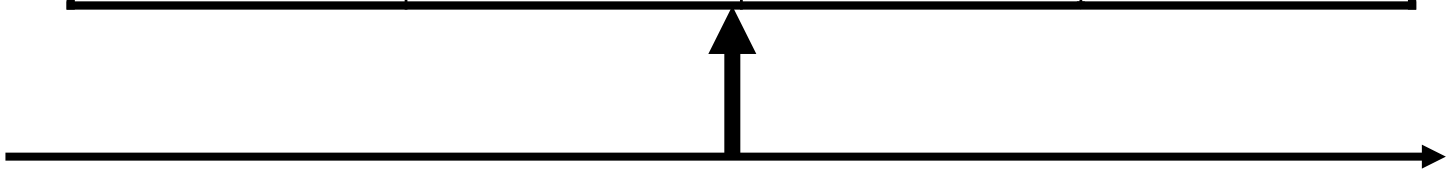

Sampling

- Systematic sampling
  - Pick out every 100th packet and record entire packet/record header
  - Ok only if no periodic component in process
- Random sampling
  - Flip a coin for every packet, sample with prob. 1/100
- Time-based sampling
  - Record a link load every n seconds
- Hash-based sampling
  - Record a packet if hash(packet) matches certain values

Inferring Results from Samples

- What can we infer from samples?
- Easy:
  - Metrics directly over variables of interest, e.g., mean, variance etc.
  - Confidence interval = "error bar"
    - decreases as $1/\sqrt{n}$
- Hard:
  - Rare events: "number of SYN packets sent from A to B"
  - Events such as: "has X received any packets"?

Stratified Sampling

- Sampling objects with different weights
  - Weight = flow size
  - Estimate average flow size
  - Problem: a small number of large flows can contribute very significantly to the estimator
- Stratified sampling: probability depends on weight
  - Sample “per byte” rather than “per flow”
  - Try not to miss the “heavy hitters”

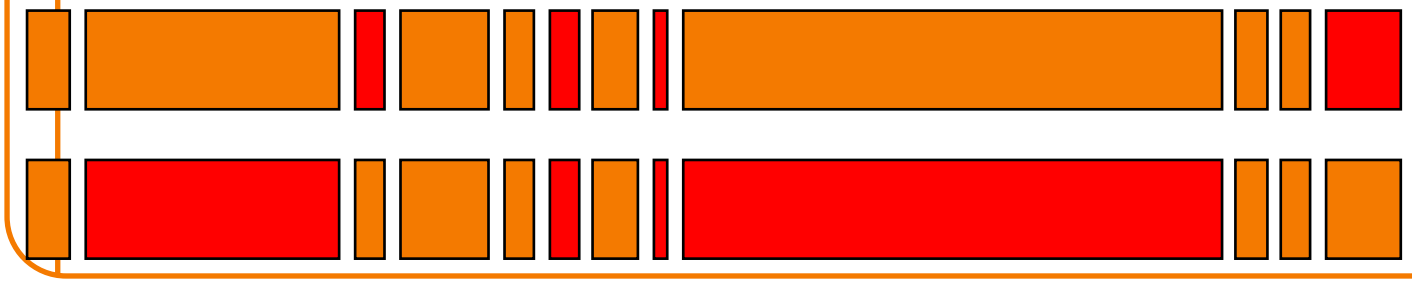

$p(x)$ constant

$p(x)$ increasing

Summary: Data Reduction Techniques

	Filtering	Aggregation	Sampling
Precision	exact	exact	approximate
Generality	constrained a-priori	constrained a-priori	general
Local Processing	filter criterion for every object	table update for every object	only sampling decision
Local memory	none	one bin per value of interest	none
Compression	depends on data	depends on data	controlled

Getting What the Operators Need

Example Challenges

- **Aggregation**: summarize the important traffic
  - Fixed level of aggregation is not sufficient
  - Want to identify large traffic aggregates
  - … at the appropriate level of aggregation
  - E.g., Jen doing an FTP, theory group web surfing
  - Example: AutoFocus, hierarchical heavy hitters
- **Sampling**: network-wide visibility
  - Managing the network as a whole
  - Coordinated measurement across different routers
  - To sample the same packets along entire path
  - … or to intentionally measure *different* traffic per hop
  - Example: trajectory sampling, cSamp

BGP Monitoring

Motivation for BGP Monitoring

- Visibility into external destinations
  - What neighboring ASes are telling you
  - How you are reaching external destinations
- Detecting anomalies
  - Increases in number of destination prefixes
  - Lost reachability or instability of some destinations
  - Route hijacking
- Input to traffic-engineering tools
  - Knowing the current routes in the network
- Workload for testing routers
  - Realistic message traces to play back to routers

BGP Monitoring: A Wish List

- Ideally: knowing what the router knows
  - All externally-learned routes
  - Before policy has modified the attributes
  - Before a single best route is picked
- How to achieve this
  - Special monitoring session on routers that tells everything they have learned
  - Packet monitoring on all links with BGP sessions
- If you can't do that, you could always do…
  - Periodic dumps of routing tables
  - BGP session to learn best route from router

Using Routers to Monitor BGP

Talk to operational routers using SNMP or telnet at command line

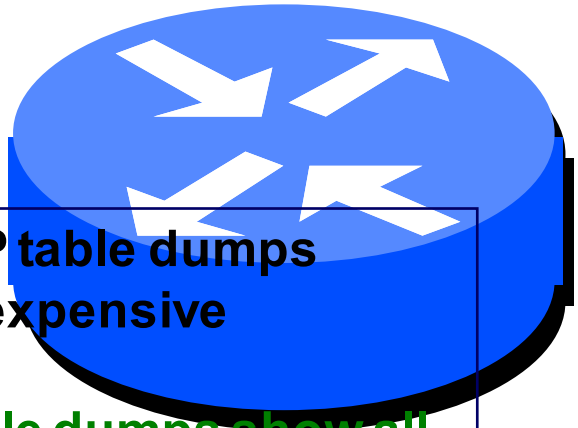

- (-) BGP table dumps are expensive
- (+) Table dumps show all alternate routes
- (-) Update dynamics lost
- (-) restricted to interfaces provided by vendors

Establish a "passive" BGP session from a workstation running BGP software

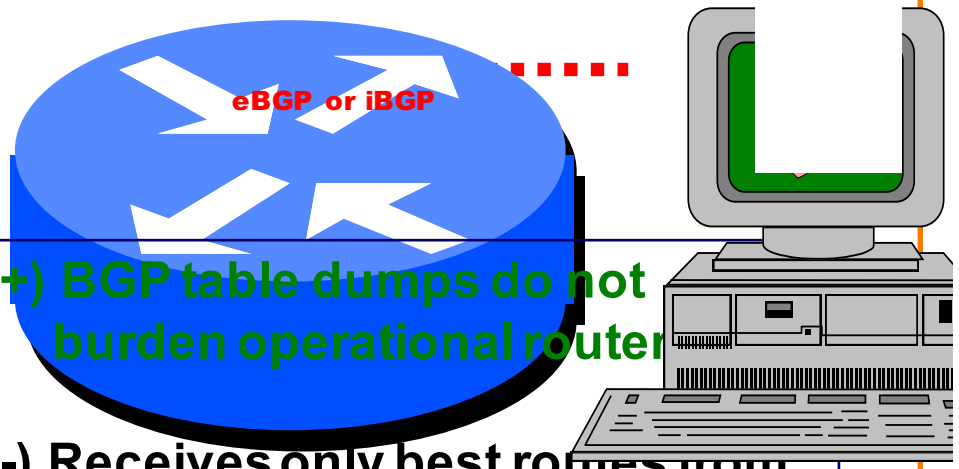

- (+) BGP table dumps do not burden operational router
- (-) Receives only best routes from BGP neighbor
- (+) Update dynamics captured
- (+) not restricted to interfaces provided by vendors

Collect BGP Data From Many Routers

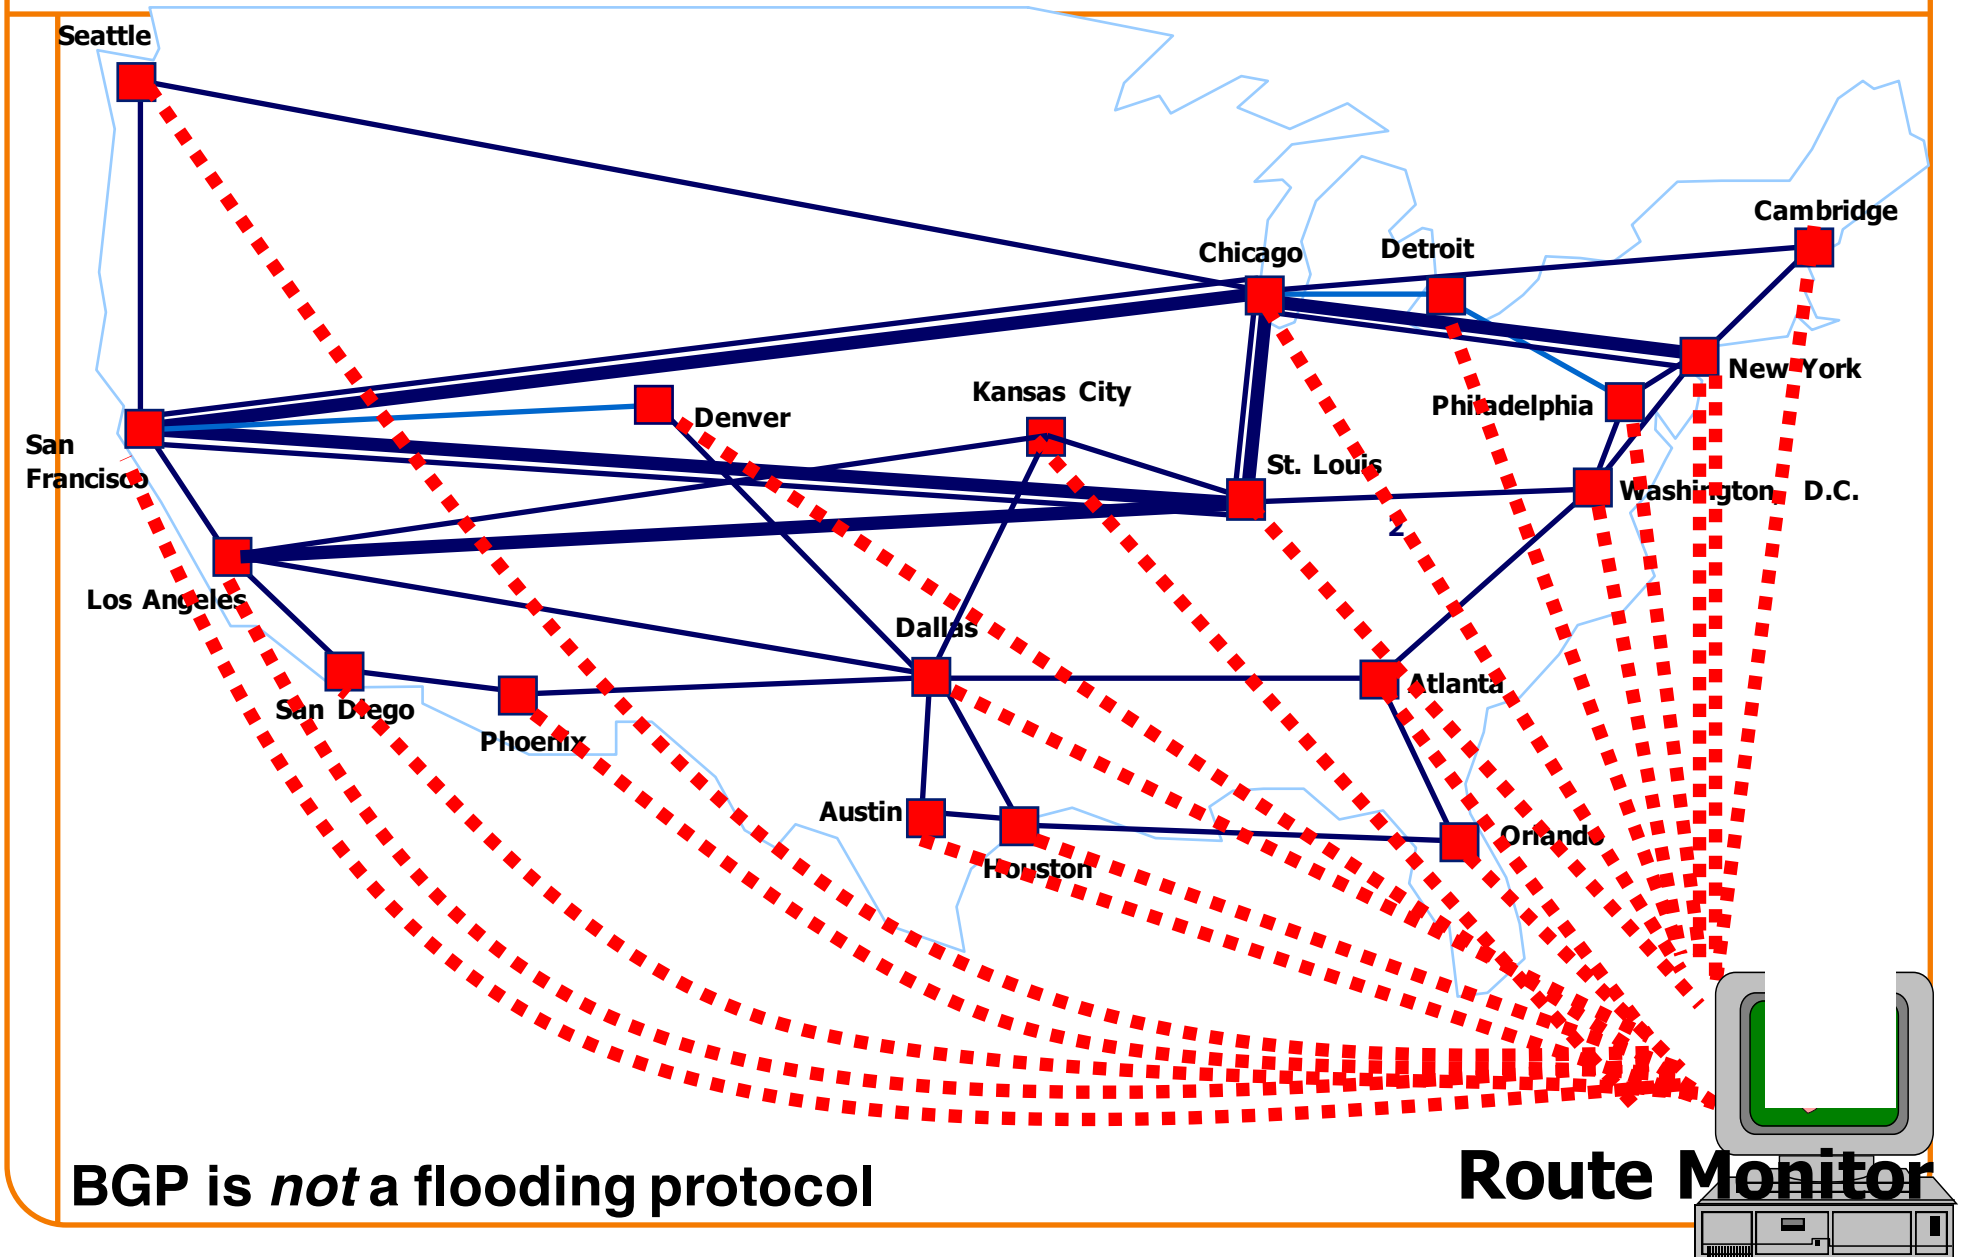

BGP is *not* a flooding protocol

BGP Table ("show ip bgp" at RouteViews)

	Network	Next Hop	Metric	LocPrf	Weight	Path
*	3.0.0.0	205.215.45.50			0	4006 701 80 i
*		167.142.3.6			0	5056 701 80 i
*		157.22.9.7			0	715 1 701 80 i
*		195.219.96.239			0	8297 6453 701 80 i
*		195.211.29.254			0	5409 6667 6427 3356 701 8
*>		12.127.0.249			0	7018 701 80 i
*		213.200.87.254	929		0	3257 701 80 i
*	9.184.112.0/20	205.215.45.50			0	4006 6461 3786 i
*		195.66.225.254			0	5459 6461 3786 i
*>		203.62.248.4			0	1221 3786 i
*		167.142.3.6			0	5056 6461 6461 3786 i
*		195.219.96.239			0	8297 6461 3786 i
*		195.211.29.254			0	5409 6461 3786 i

AS 80 is General Electric, AS 701 is UUNET, AS 7018 is AT&T
AS 3786 is DACOM (Korea), AS 1221 is Telstra

Conclusions

- Measurement is crucial to network operations
  - Measure, model, control
  - Detect, diagnose, fix
- Network measurement is challenges
  - Large volume of measurement data
  - Multi-dimensional data
  - Multiple data sets
- Great space of research problems
  - Streaming algorithms, statistical inference, …
  - New hardware-efficient measurement primitives
  - Analysis of existing measurement data
  - …