Resource Management in an Active Measurement Service

Ethan Blanton[†], Sonia Fahmy[†], Sujata Banerjee[‡]

[†]Purdue University [‡]Hewlett-Packard Labs

April 18, 2008

The Setting

Many applications use network measurement or knowledge of network characteristics to operate, or to enhance operation.

- Azureus
- Tor
- End-System Multicast
- Experimental overlays

Future network services seem likely to continue this.

The Setting (cont'd)

Measurement *as a service* has several potential benefits.

- Elimination of duplication
- Measurement techniques lifted out of applications
- Control and accounting of resource consumption
- …

The Setting (cont’d)

Measurement as a service has some issues, as well!

- Potential for abuse
  - Administratively-defined resource bounds
- Limited transparency
  - Report errors and rejected requests to the requestor
- Effects of constrained vantage points
  - Scatter measurement hosts far and wide
  - Account for differences between infrastructure and end host viewpoints

Architecture

Our system:

- Is *open*, requiring no *a priori* enrollment or identity establishment;
- Is *transparent*, communicating decisions and known conditions which might affect output to requestors;
- Is *dynamic*, scheduling measurements on demand;
- Is *simple*, requiring applications to have little knowledge of measurement tools or techniques.

This work

The focus of this work is on limiting resource consumption.

- To allow administrators to specify the allowable impact on their networks
- To prevent usage of the service as a DoS tool

We discuss:

- A *load invariant* for measurement hosts
  - *e.g.*, “Outgoing measurement traffic on the access link to host M1 shall not exceed 128 kbps”
- Simple techniques to preserve this invariant

This work (cont'd)

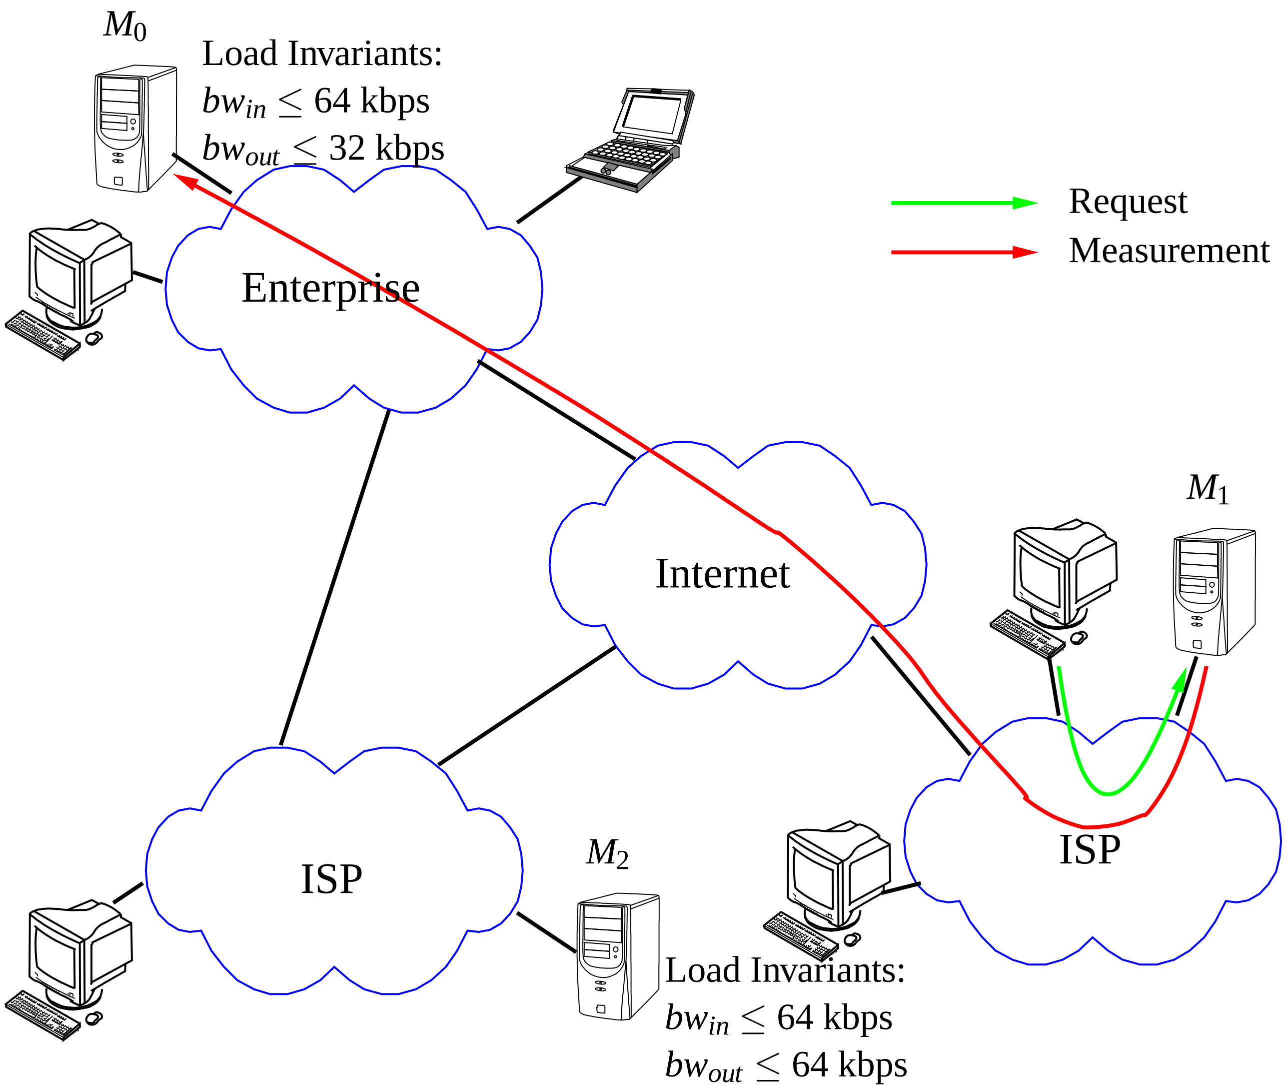

Admitting Active Measurements

Measurement requests are for periodic active measurements.

Admission is performed by:

- Comparing the requested measurement load over discrete intervals of time with an administratively defined load invariant, and
- Rejecting measurements which violate that invariant.

Measurement resource consumption is estimated:

- By a *cost vector* approximating resource requirements, which is
- Based on empirical measurements.

Assumptions

1. End hosts using the serviceneed not be trusted, but measurement hosts are.
2. Time is encoded as an absolute time in UTC, and clock drift among measurement nodes is small.
3. Measurement requests, once accepted, cannot be preempted.
4. Load invariants are maintained only for *sender* and *receiver* nodes of active measurement traffic, and their access link to the Internet.

Admission Control Tests

For now, we deal only with inbound & outbound bandwidth, in two alternative methods.

Peak Bandwidth: The maximum bandwidth required by a measurement tool in *any* interval of a specified length.

Average Bandwidth: The average bandwidth required by a measurement tool over *all* (discretized) periods of a specified length.

Prototype and Experimental Evaluation

Experiments were performed on emulab the utah network testbed (*http://www.emulab.net/*).

We aim to:

- Quantify the benefit of applying admission control
- Compare different admission control tests and timescales
- Demonstrate that our simple estimates are robust

Experimental Scenario

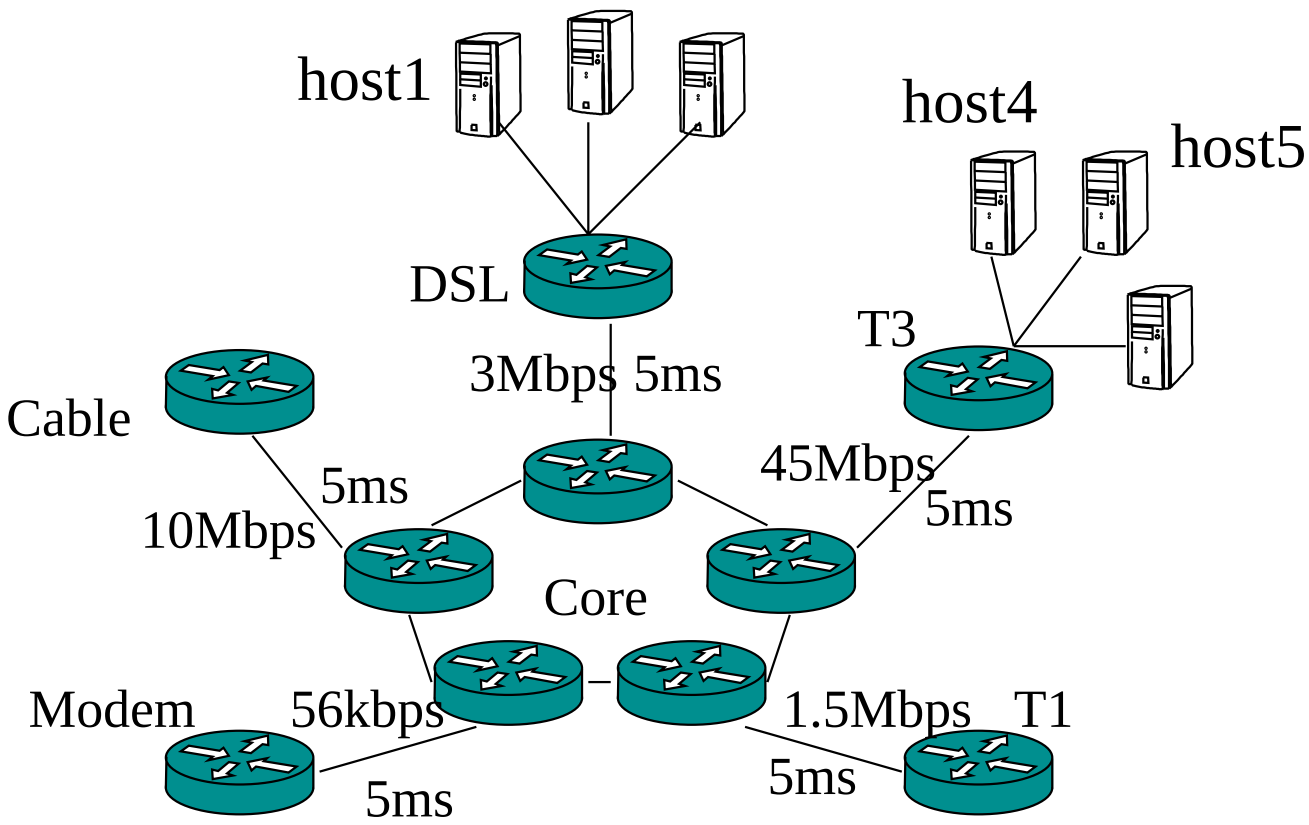

Experimental Scenario (cont'd)

- Resource limits are 20% of access link bandwidth
- Time interval for estimation is 1s
- Measurement tools are popular examples such as `ping`, `pathrate`, `tulip`...
- Measurement tool cost vectors were created with all-pairs probes
  - In practice, each node would use a ballpark estimate for its approximate bandwidth

The Case for Admission Control

Host	requested	failed	bw_{in}	bw_{out}	V_{in}	V_{out}
No Admission Control						
T3	13.2	0	441	359	0	0
DSL	11.7	0	199	451	345	838
Modem	11.8	0	15	20	602	1151
Cable	13.3	0	446	209	0	0
Average Admission Control						
T3	13.2	3	278	295	0	0
DSL	11.7	4	129	247	130	415
Modem	11.8	6.6	4	4	223	224
Cable	13.3	3.6	311	154	0	0
Peak Admission Control						
T3	13.2	5.1	113	211	0	0
DSL	11.7	5.7	31	57	0	0
Modem	11.8	7.8	3	3	78	56
Cable	13.3	5	217	91	0	0

Estimation Accuracy

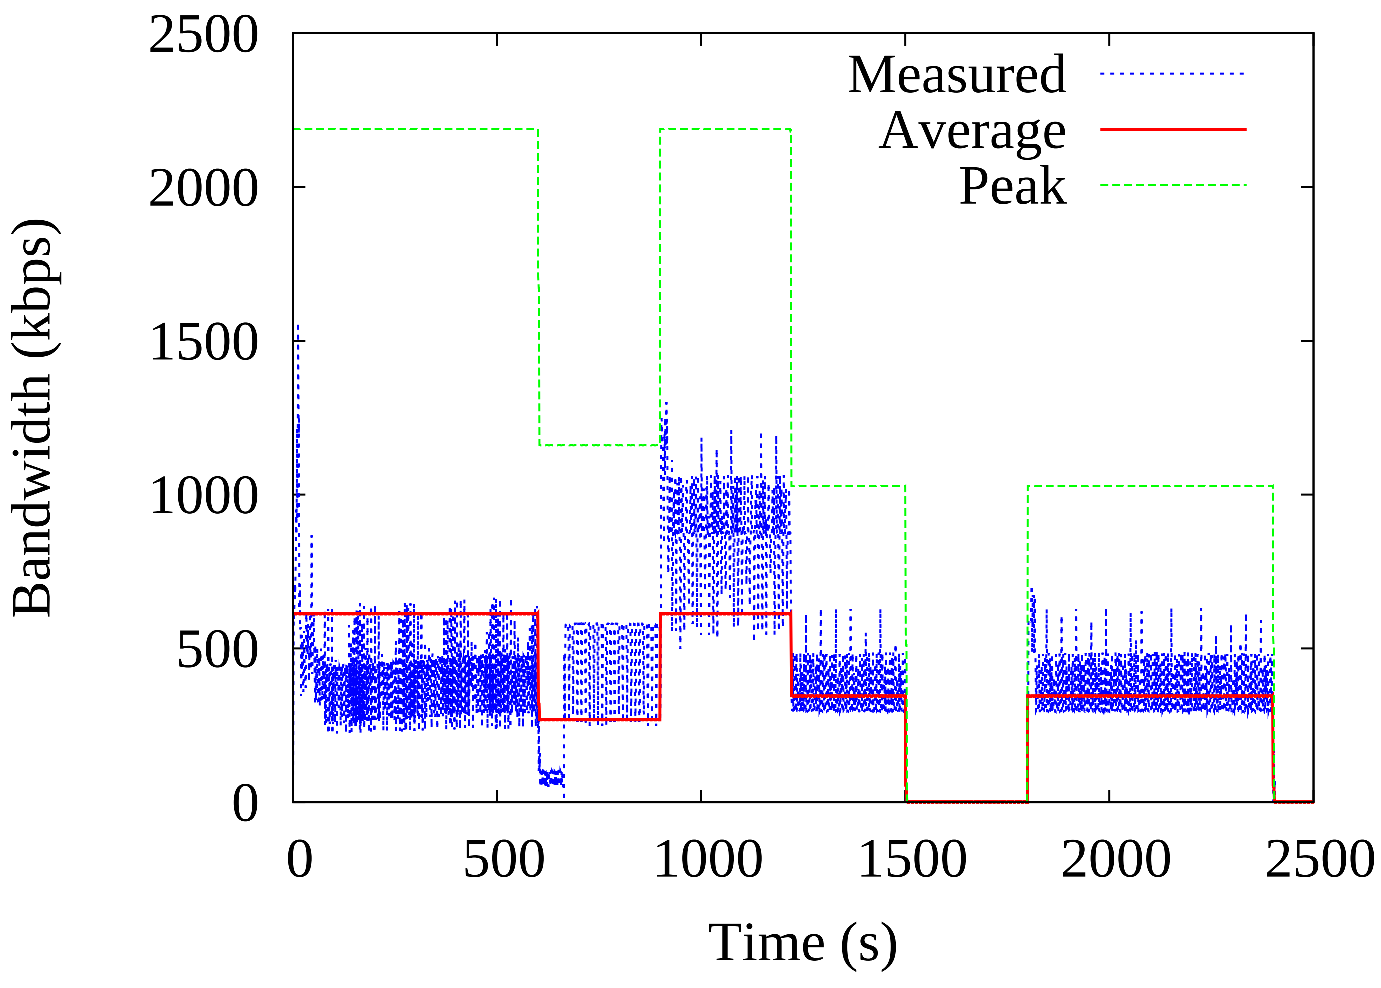

Take-away

- Using global averages more closely approximates aggregate bandwidth utilization, but still misses some detail.
- Pessimistically assuming that resource utilization peaks will coincide is overkill, at least for these workloads.
- Average bandwidth results in some under-estimation, leading to potential violations.

Changing Network Conditions

- Networks are not static, and paths vary widely in character.
- We know that end-to-end bandwidth affects the bandwidth utilization of some measurement tools (*e.g.*, some bandwidth capacity estimation tools) significantly.
- We look at the impact of unpredictable cross-traffic and changes in end-to-end delay.
  - Cross-traffic: small, tolerable perturbation in measurement tool bandwidth utilization.
  - End-to-end delay: No correlation between delay changes and measurement tool bandwidth utilization.

Estimation Timescales

The measurement tool estimation process estimates measurement tool cost vectors over some timescale.

The absolute value of this timescale does not seem to be critical (we evaluated 100 ms, 1 s, 2 s, and 4 s).

- Average bandwidth consumption is unaffected by timescale.
- Peak bandwith consumption changes, but its pessimistic nature limits its scheduling effects to an acceptance rate change of about 10% in either direction.

Conclusions

- Estimation of measurement tool cost vectors is important for *a priori* measurement scheduling.
- Peak bandwidth utilization as a scheduling metric is overly pessimistic due to the bursty nature of measurement tool traffic.
- Average bandwidth utilization captures some of the bursty behavior when measurements can be multiplexed.
- While these simple methods show promise and feasibility, more complex models to capture measurement tool cost vectors may allow more accurate admission decisions.

Questions?

Related Work

[1] E. Blanton, S. Fahmy, and S. Banerjee. A framework for an on-demand measurement service. Technical report, Purdue University, 2008. http://www.cs.purdue.edu/homes/fahmy/reports/measurement.pdf.

[2] P. Calyam, C.-G. Lee, E. Ekici, M. Haffner, and N. Howes. Orchestration of network-wide active measurements for supporting distributed computing applications. *IEEE Transactions on Computers*, 56(12), Dec 2007.

[3] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: An information plane for distributed services. In *Proc. of OSDI*, pages 367–380, November 2006.

[4] A. Nakao, L. Peterson, and A. Bavier. A routing underlay for overlay networks. In *Proc. of SIGCOMM*, pages 11–18, 2003.

[5] N. Spring, D. Wetherall, and T. Anderson. Scriptroute: A public internet measurement facility. In *Proc. of USITS*, 2002.

[6] P. Yalagandula, P. Sharma, S. Banerjee, S. Basu, and S.-J. Lee. s^3: A scalable sensing service for monitoring large networked systems. In *Proc. of INM*, September 2006.