First PlanetLab Asia Workshop

Happy Packets to You!

Randy Bush
Timothy G. Griffin
Jun Li
Z. Morley Mao
Eric Purpus
Dan Stutsbach

This research is funded by NSF award ANI-0221435.

Jun Li, U of O, 09/17/04

Outline

- Why this study?
- Methodology
- Results and analysis
- Open issues
- Conclusions

Why This Study?

- We frequently hear comments about Internet control plane quality, such as
  - Internet routing is fragile and collapsing,
  - Yesterday was a bad routing day on the Internet,
  - BGP is broken or is not working well,
  - Changing protocol X to Y will improve routing, or
  - Internet routing has been severely affected by event X (e.g. power blackout, worm outbreak)
- But what measurement can really tell the quality of control plane?
  - Number and frequency of BGP updates?

Happy Packets

- What ultimately counts is whether the customer's packets can reach their intended destination with good performance
  - Namely, the performance at data plane
  - And after all, this is the functionality of the control plane
- We call them happy packets

Our Goal

- Answer this question: Are packets happy under routing changes?
  - Basically, we evaluate Internet control plane quality by measuring the data plane performance

Methodology

- How to measure packet happiness at the data plane?
  - Use the PlanetLab
- How to introduce routing changes into the control plane
  - Use a BGP Beacon

Happiness over the PlanetLab

- A set of geographically and topologically diverse PlanetLab nodes are selected as probe sites
- A site from Seattle is selected as a sink
- Every probe site continuously sends testing UDP streams toward the sink site
  - While the routing toward the sink changes
- Over a period of four months

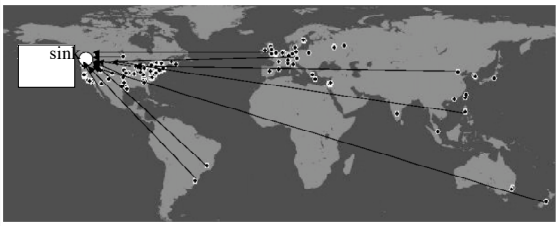

Metrics of Happiness

- Using well-established metrics (DDJ&R):
  - Delay, drop, jitter, and reordering
- Delay: the relative to the mean one-way delay
- Loss rate: % of packets dropped per second
- Loss duration: the length of a time window with exceptionally high loss rate
- Jitter: delta between delays
- Reordering rate: % of reordered packets per second

BGP Beacon

- An unused globally visible IP address prefix
- With a schedule of BGP announcements and withdrawals regarding reaching the prefix
- We use a multi-homed BGP Beacon 192.83.230.0/24
- The test stream sink has a specific IP address from this prefix

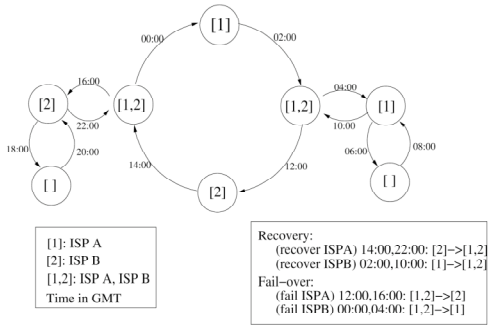

Collecting Control Plane Data

- Oregon RouteViews Project archive BGP updates
- Can help observe BGP updates related to the BGP Beacon prefix
  - Thus BGP duration and BGP update number during an event can be measured

What Did We Find?

- Average and worst case DDJ&R results of individual streams
- Aggregated results of DDJ&R
- Control plane data

Results of an Average Stream

- Stream from 128.95.219.192 as an example
- Under four different routing changes over 20 min: AB-B, AB-A, A-AB, B-AB
- Performed well in general, either during or outside routing changes
  - Most times packet delays are acceptable
  - No reordering was detected
  - Thus jitter is also acceptable
  - A 30-sec loss duration in the AB-A case

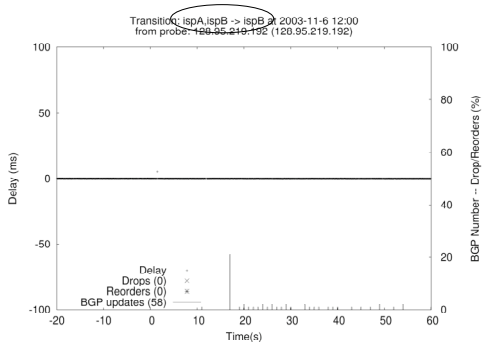

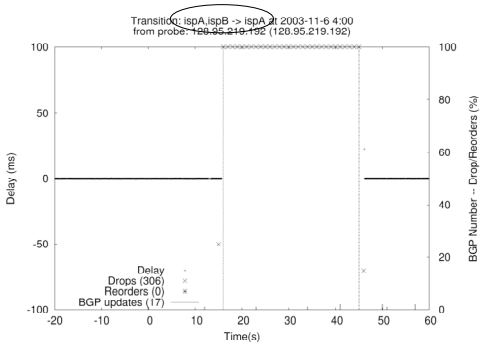

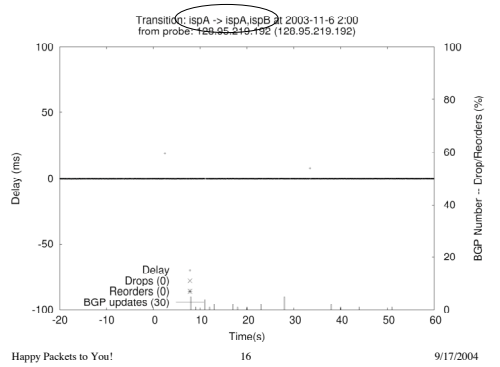

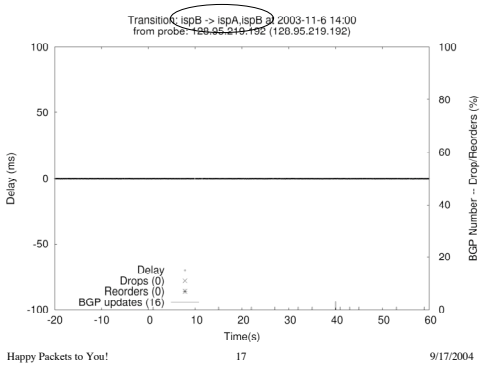

Results of a Worst-Case Stream

- Stream from lcs-bgp.vineyard.net
- Under four different routing changes over 20 min: AB-B, AB-A, A-AB, B-AB
- Performed the worst compared to other streams
  - Longer delay than others
  - Longest loss duration
    - 10s in AB-B w/ 91 drops & 8 reorders
- However, not significantly worse than its own normal period

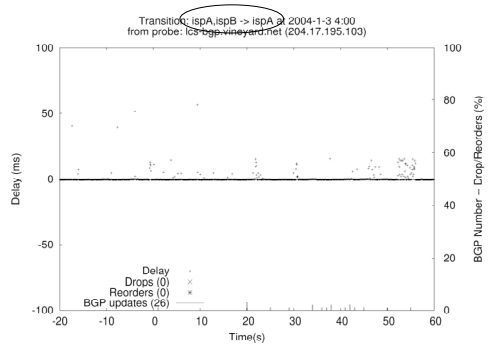

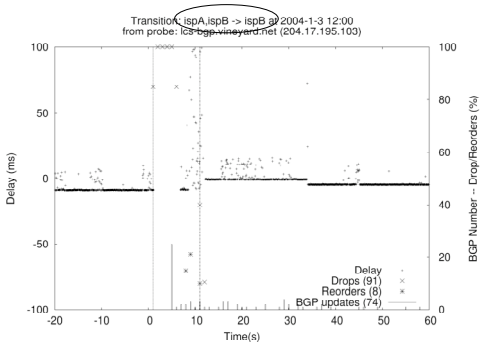

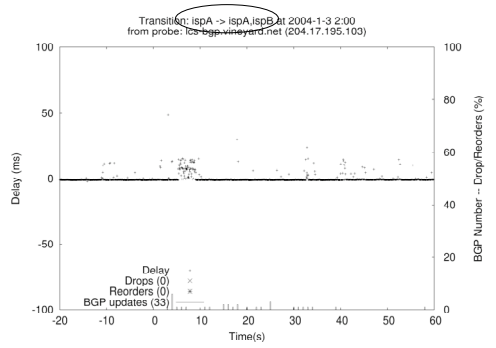

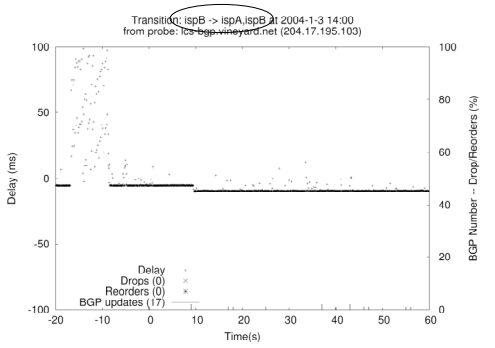

Aggregated Results

- Delay CDF
- Jitter CDF
- Loss rate
- Reordering
  - not plotted (close to 0)

CDF of (Relative) Delay

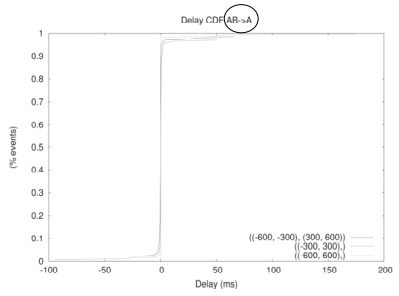

Results for AB-B, A-AB, B-AB are similar

Jitter CDF

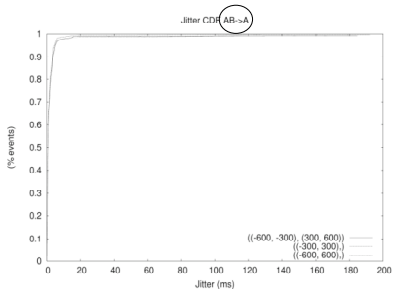

Results for AB-B, A-AB, B-AB are similar

Loss Rate

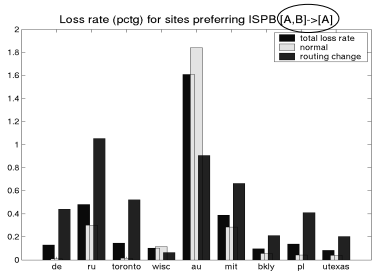

Loss Rate

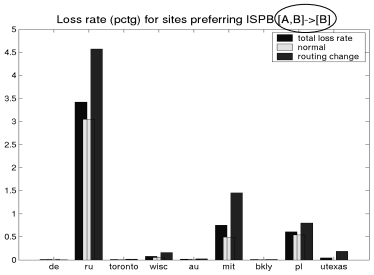

DDJ&R Summary

- Acceptable during injected routing changes
  - Although generally worse than normal periods
- In most cases, BGP performs well
- Can also approximate closely the packet delivery performance between two routers
  - DDJ&R of each UDP stream from end to end is also close to the BGP Beacon router and the router for the probe site

Misconception in Inferring Packet Happiness

- Can control plane data, such as those from RouterViews or RIPE, predict packet performance, thus equivalent to measuring DDJ&R?
- Answer: No!

Loss duration vs. BGP update duration

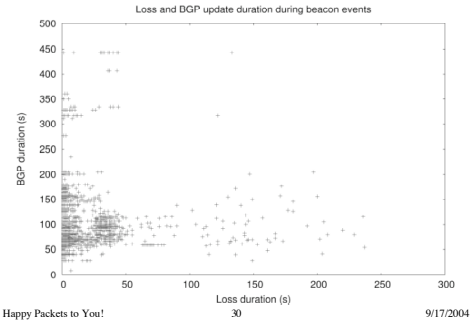

Loss duration vs. BGP update number

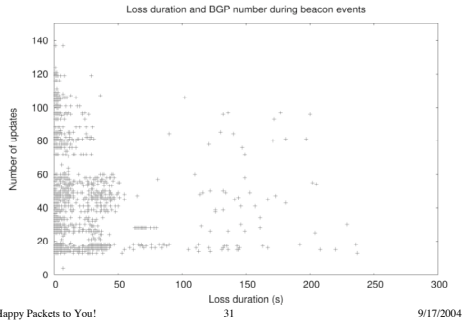

White Blood Cells

- Perhaps BGP announcements are like white blood cells
- Their presence may signal a problem
- But they are often part of the cure, not necessarily part of the problem

Summary

- No clear correlation between loss duration and BGP duration, or loss duration and number of BGP updates
- RouteViews archives only provide partial knowledge of the control plane
- One should be cautious in using BGP updates to analyze control plane quality

Open Issues

- Large-scale control plane events
- Congestion effects on DDJ&R
- Usage of partial control plane knowledge
- . . . .

Conclusions

- Data plane performance is the best measure of control plane effectiveness
- Not only for BGP, but also other routing protocols
- We've found little proof that BGP is not resilient or performing poorly during routing changes
- And we should be critical about using partial control plane data for study

Questions?

Randy Bush <randy@psg.com>
Timothy G. Griffin <tim.griffin@intel.com>
Jun Li <lijun@cs.uoregon.edu>
Z. Morley Mao <zmao@eecs.umich.edu>
Eric Purpus <epurpus@cs.uoregon.edu>
Dan Stutsbach <agthorr@cs.uoregon.edu>