Happy Packets
a Second Experiment

APNIC / Kyoto

2005.02.23

Randy Bush <randy@psg.com>

Matt Roughan <matthew.roughan@adelaide.edu.au>

<http://rip.psg.com/~randy/050223.happy2-apnic.pdf>

Central Question

- What is the relationship between **control plane** instability and **data plane** instability?

- Related Questions:
  - Is the quantity of BGP updates good or bad?
  - Who wants to see zero BGP updates?

Internet Weather

We frequently hear comments such as

- Internet routing is fragile, collapsing, ...,
- BGP is broken or is not working well,
- Day X was a bad routing day on the internet,
- Change X to protocol Y will improve routing,
- Etc.

And we often measure routing dynamics and say that some measurement is better or worse than another

Internet [Routing] Instability

- We are told that a lot of BGP updates is equated with internet instability
- "There are too many BGP updates, so BGP must be broken."

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

Routing Quality

- But what is *good* routing? How can we say one measurement shows routing is better than another unless we have **metrics** for routing quality?
- We often work on the assumption that number of prefixes, speed or completeness of convergence, etc. are measures of routing quality

Happy Packets

- The measure which counts is whether the users' packets reach their destination
- If the users' packets are happy, the routing system, and other components, are doing their job
- We call these ***Happy Packets***
- There are well-known metrics for the data plane, Delay, Drop, Jitter, and Reordering
- So we set out to measure Control Plane quality by measuring the Data Plane

Experiment One

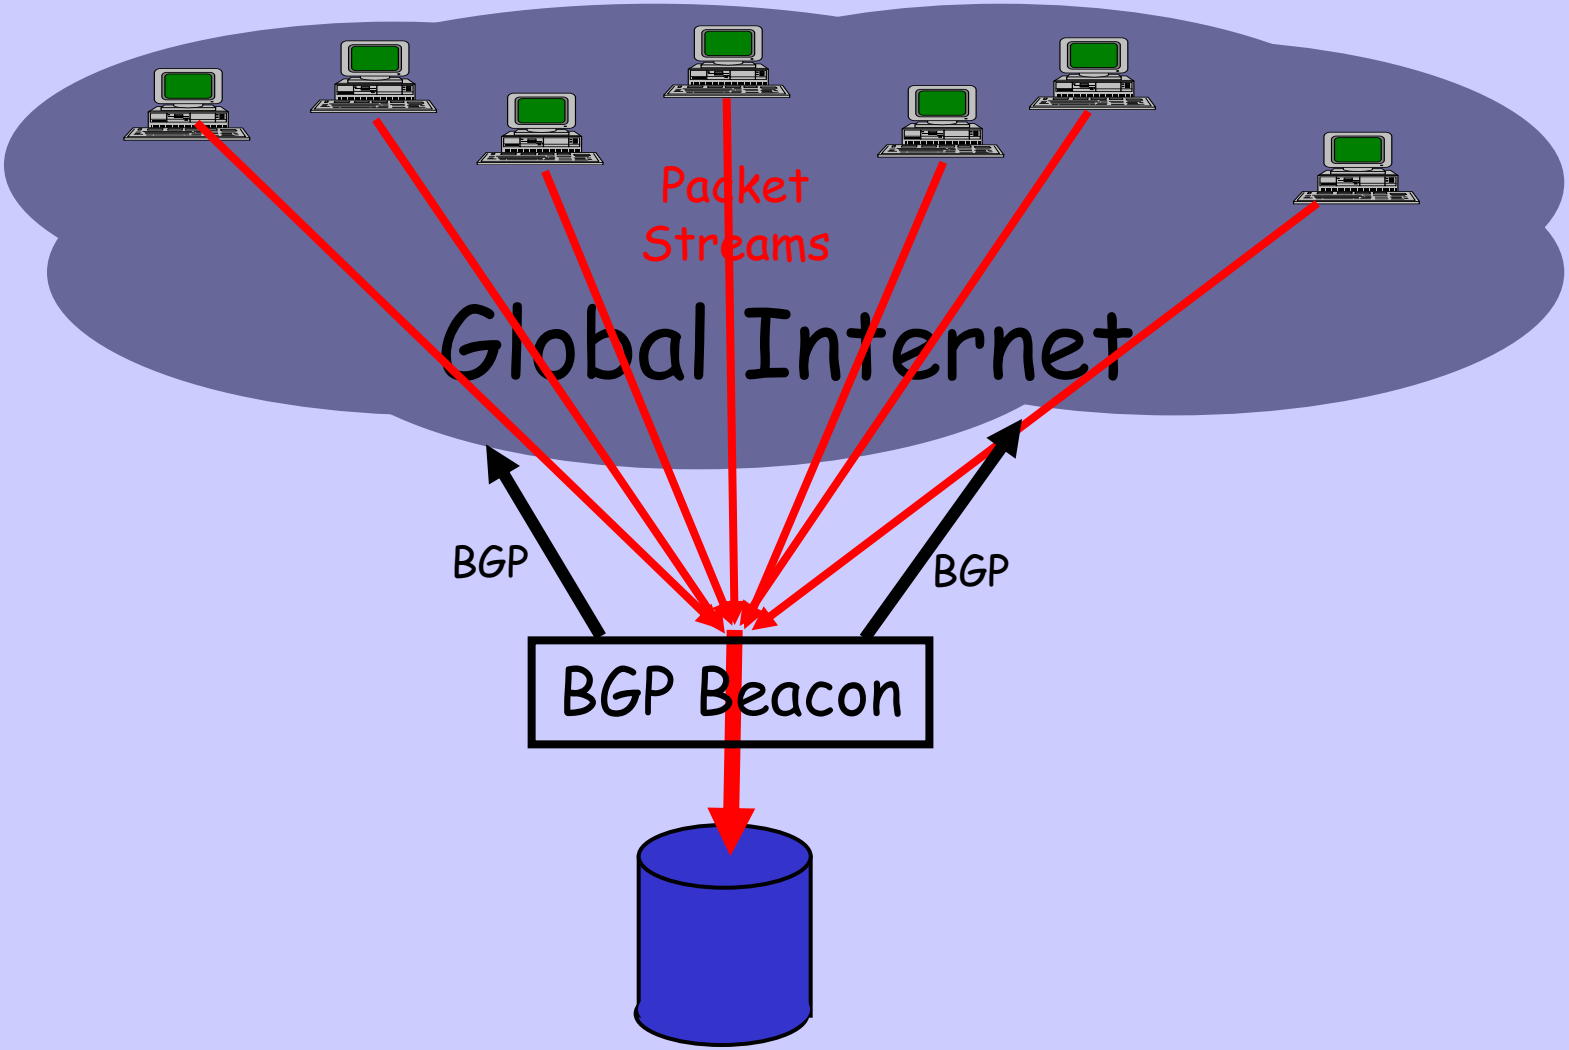

Experiment One

- Artificial injection of routing updates and measured packet performance toward the routing target
- Found **no** significant correlation between number or time of updates and data performance
- But this was artificial and did not test for large scale real events

Experiment Two

- So, we looked at some large internet events: Code Red, Nimbda, and Slammer
- Route-Views gave us the control plane, the BGP announcements
- RIPE TTM Project gave us the data plane, packet performance data

RIPE TTM Boxes

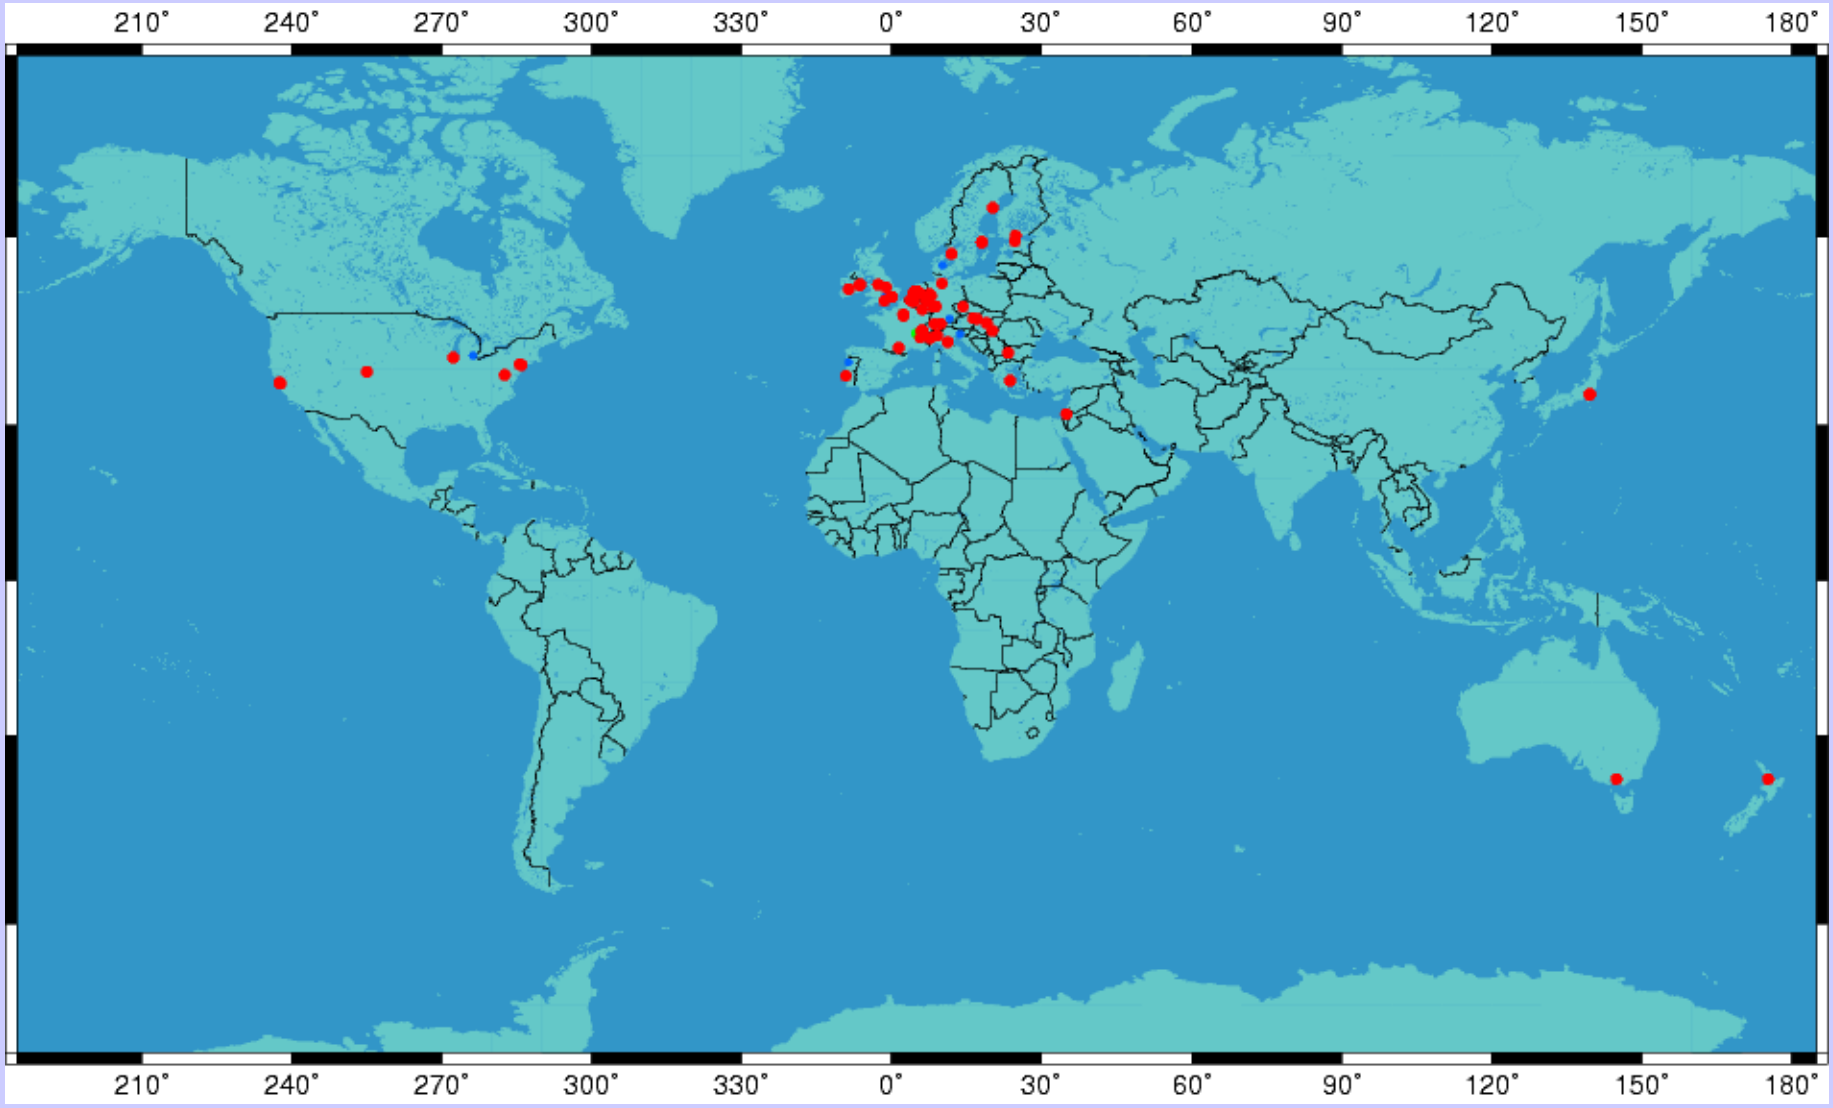

Code Red – Delays & BGP Counts

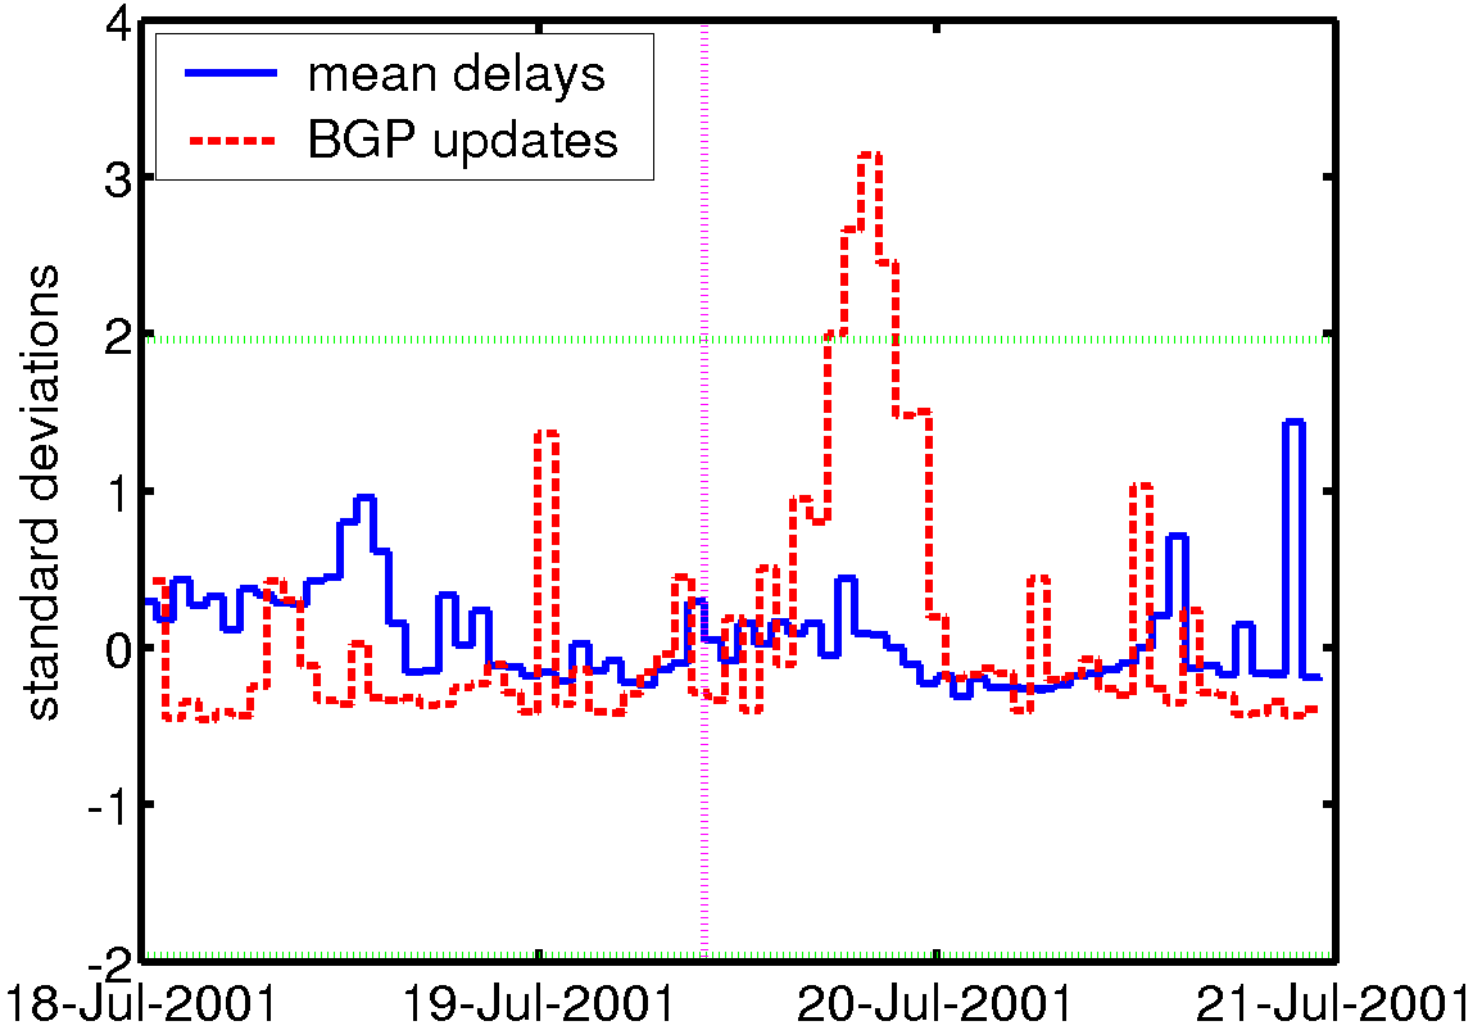

Nimbda – Delay & Updates

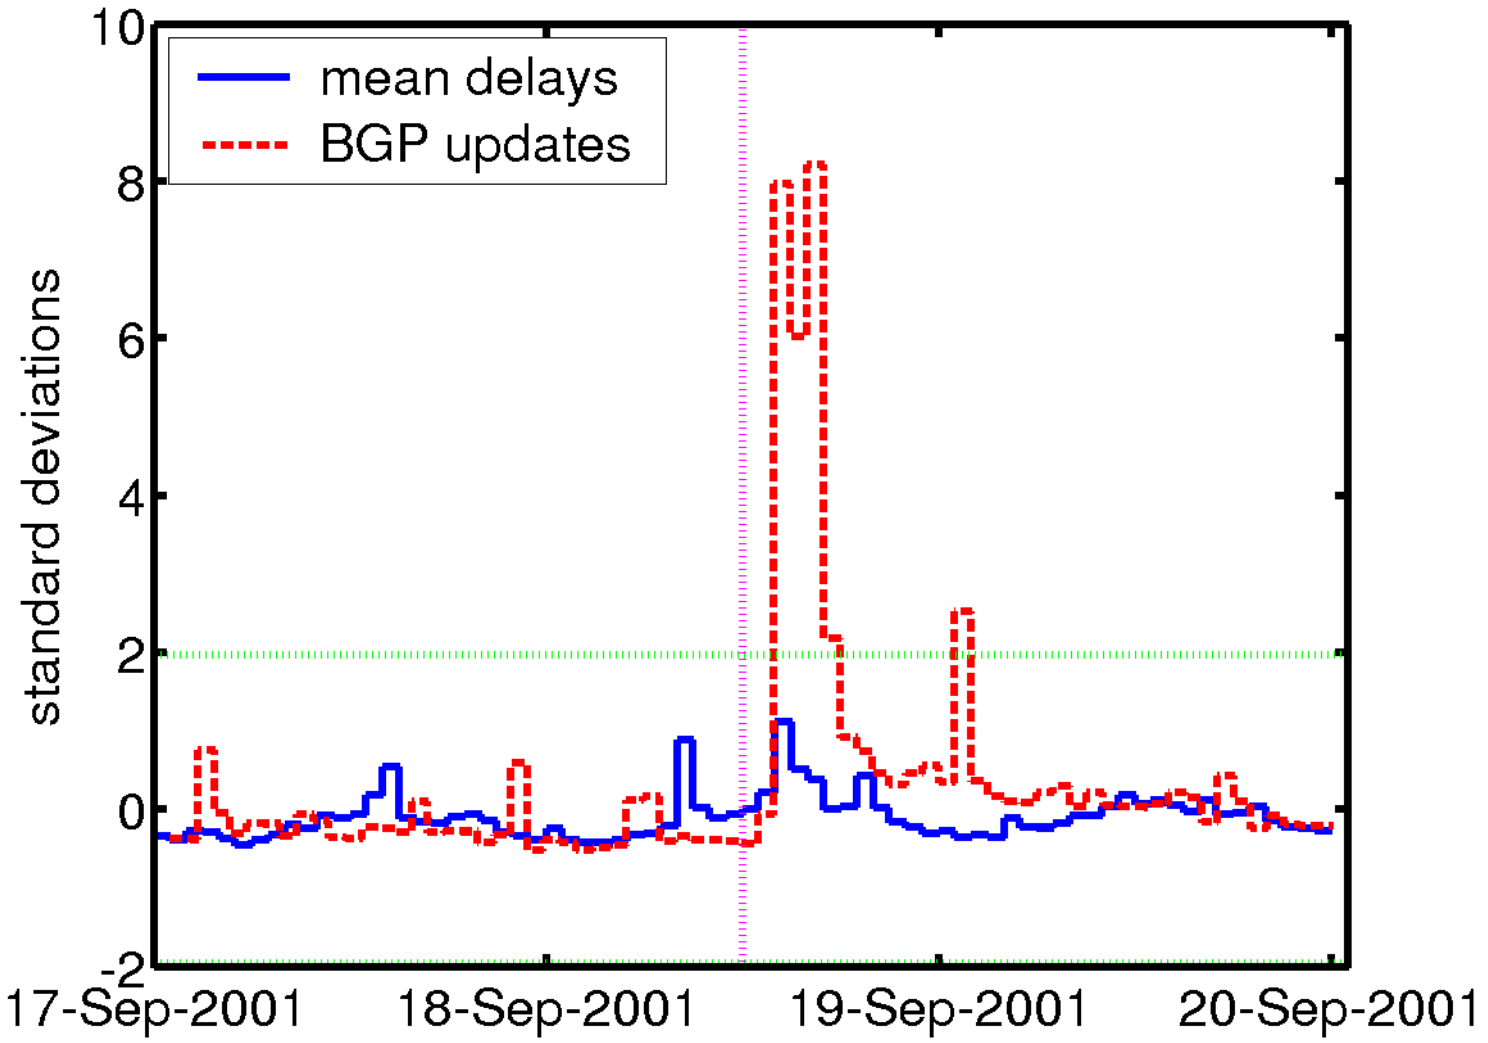

Slammer – Delay & Updates

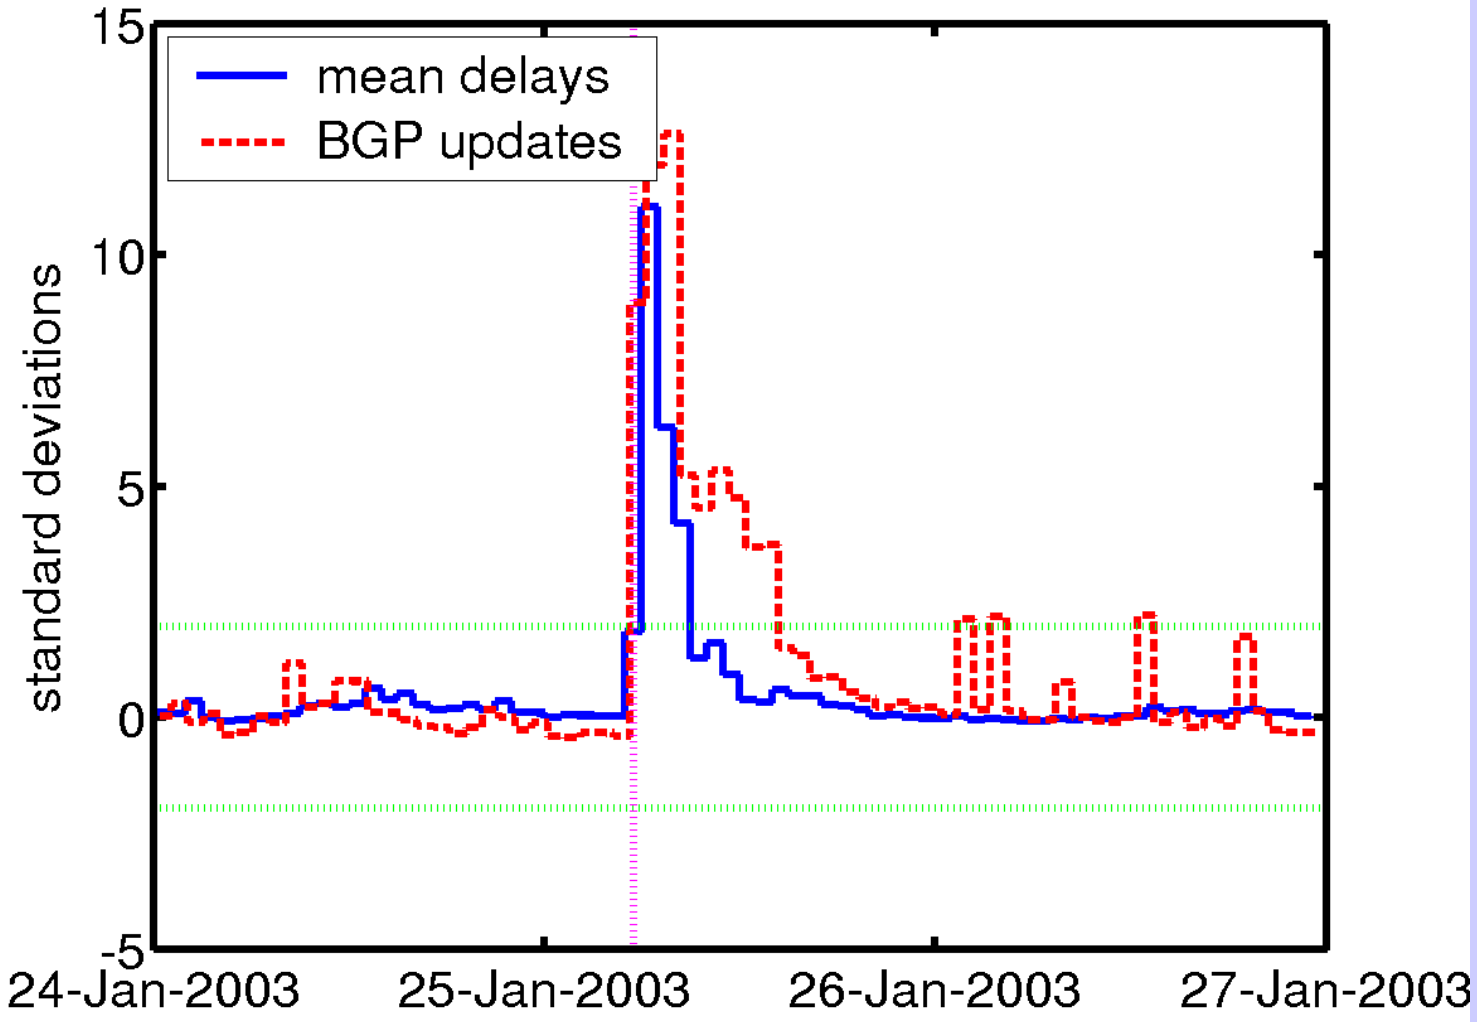

Time Series – Red & Nimbda

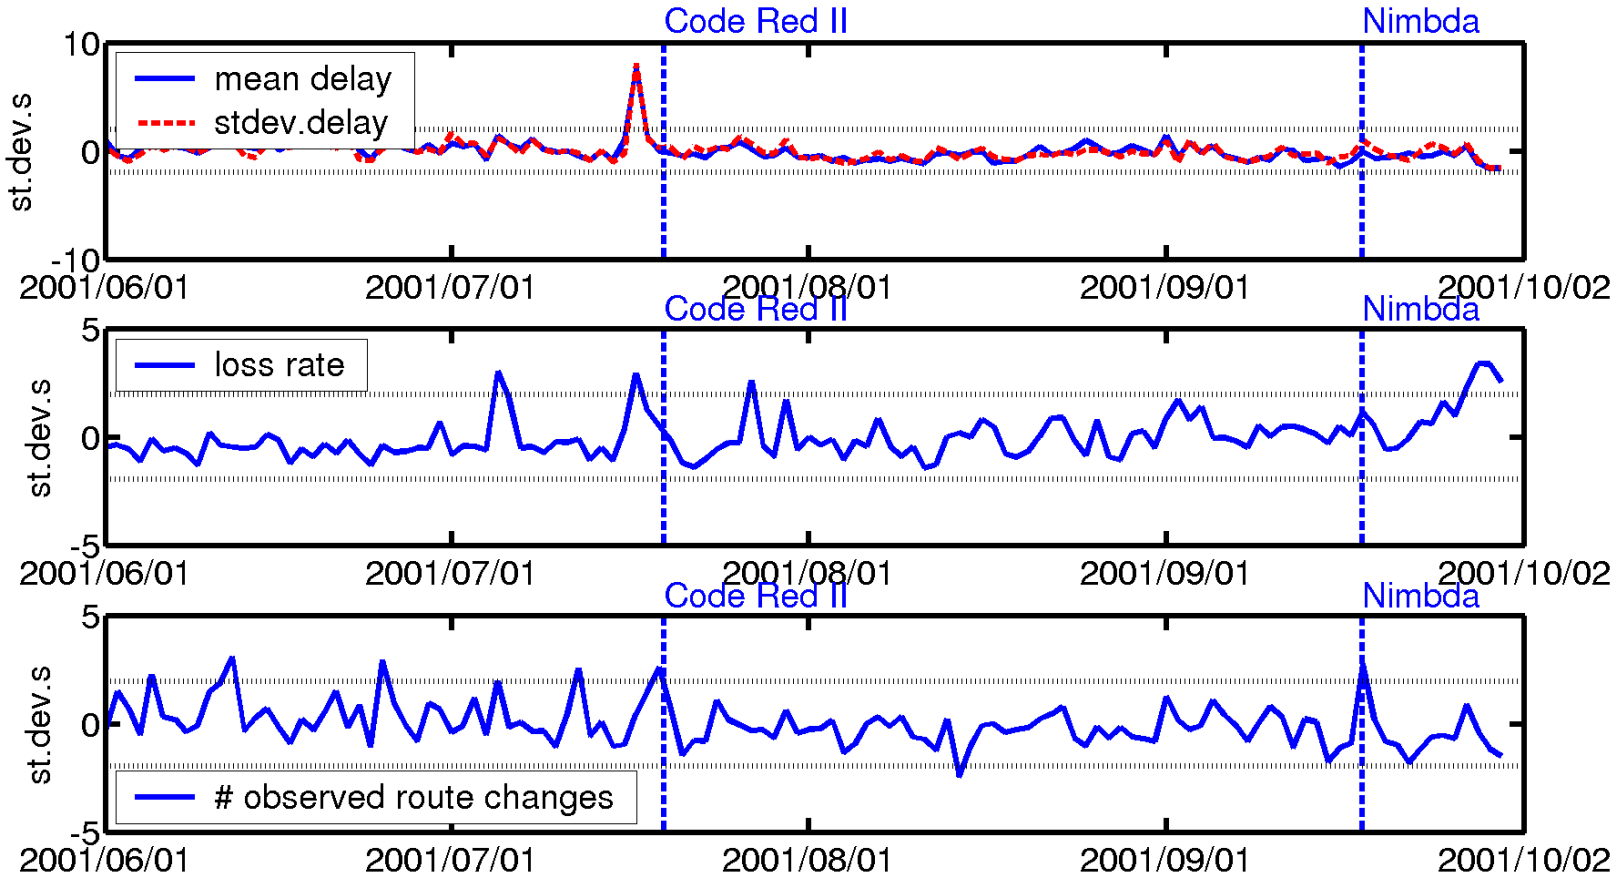

Time Series - Slammer

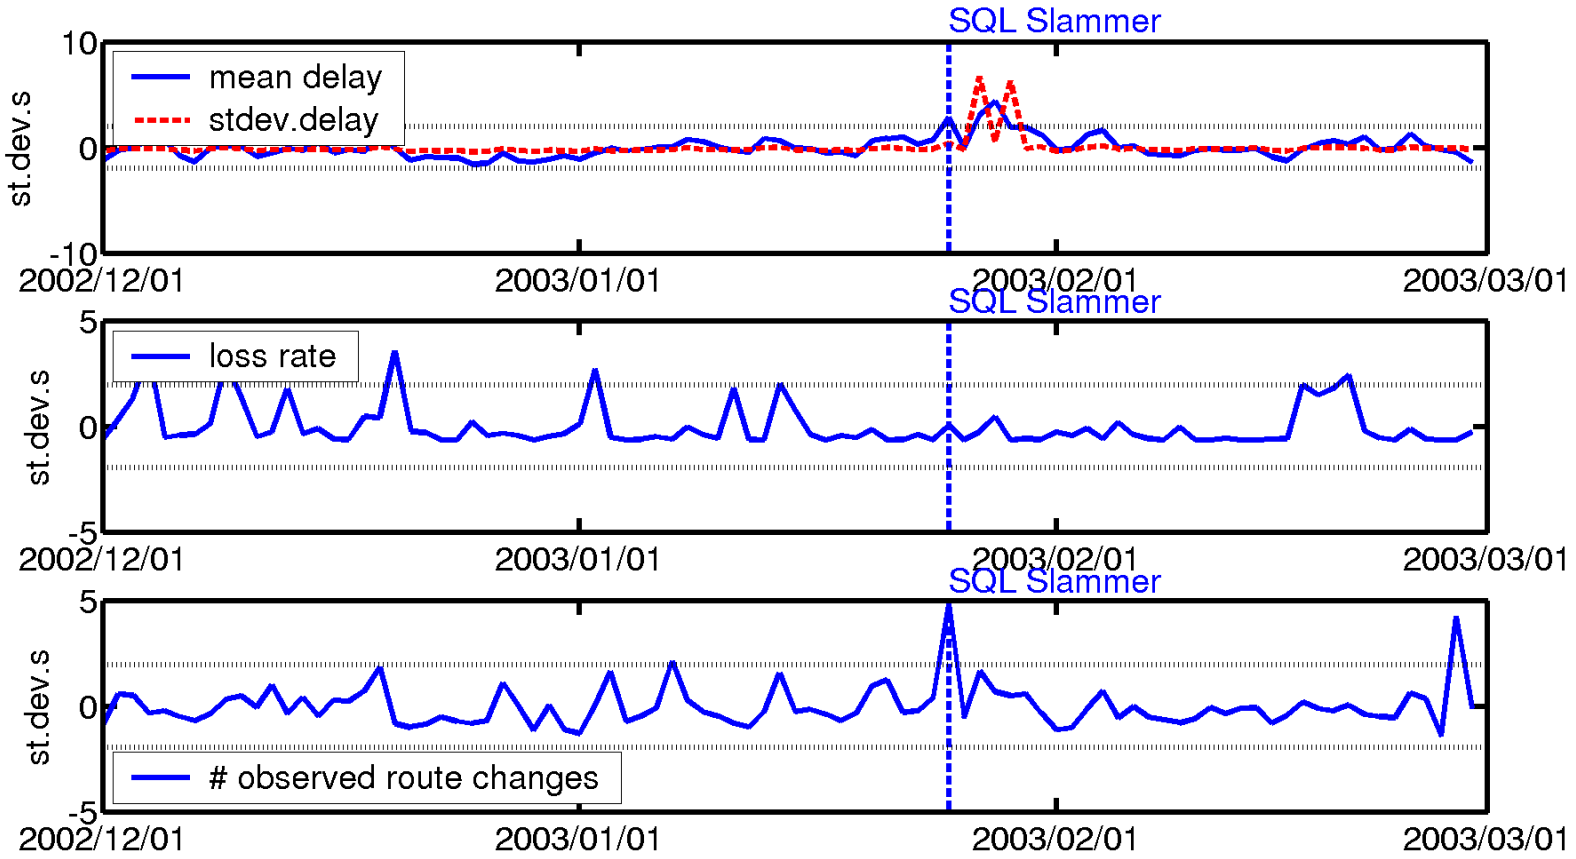

Thoughts

- Watching BGP update count or frequency, though easy, is **not** a good predictor of user experience
- Measure performance **directly**
- Delay, Drop, Jitter, & Reordering are well-known and measurable, use them
- Would be nice to have more RIPE TTM boxes in Asia/Pacific

Thanks To

- NSF via award ANI-0221435 (Oregon)
- ARC grant DP0557066 (Adelaide)
- The University of Oregon Route-Views (Lucy and Joel)
- RIPE/NCC Test Traffic Project
- The University of Adelaide (Matt)
- Internet Initiative Japan (Randy)
- Verio and Sprint (bandwidth)
- Juniper, Cisco, & Procket (routers)