Real-time Global Routing Metrics

A Common Language for Measuring Routing Stability

Jim Cowie, Andy Ogielski, BJ Premore, Eric Smith, Todd Underwood

Renesys Corporation

Why care about metrics?

1. Transient withdrawals and flapping routes affect performance on all our networks, whoever causes it.

2. Operators need to be able to evaluate the scope and magnitude of routing instabilities in their own networks and determine the appropriate level of response.

3. Metrics offer the community a common language to objectively describe and compare the impact of network events.

4. We cannot improve what we don’t measure.

In this presentation, we:

1. Introduce fast-changing routing metrics for misbehaving prefixes – prefixes for which most routes are flapping or transiently unavailable.

2. Propose two concrete fast metrics with intuitive meaning (NetsOut and PenaltyBoxes) and illustrate their behavior.

3. Show how to compute them using your own routing data, or UPDATE traces from RouteViews, RIPE, or other BGP archives.

Two kinds of routing metrics

"Slow structural metrics" **merge routing tables or UPDATEs to study evolution of structure: AS interconnectivity graph, path lengths, prefix statistics.**

...Much good work already done: Broido, Chang, claffy, Gao, Govindan, Huston, Jamin, Rexford, Shenker, Willinger, and many others.

"Fast instability metrics" **correlate UPDATEs from diverse BGP sessions to characterize evolution of path dynamics: transient outages, flapping routes,...**

...Simple counting metrics (routing table sizes, various UPDATE rates,...) are not sufficiently informative.

An ideal instability metric should be…

- Equally applicable to prefixes in a single AS, single country, or the whole Internet.
- Meaningful across multiple timescales, from minutes to years.
- Tunable, but insensitive to computational details.
- Easy to understand and easy to compute.

…Our goal here: **Start Simple, Stay Empirical**.

Simple Fast Instability Metrics

NetsOut: ***count the subpopulation of transiently withdrawn network prefixes at any given moment.***

PenaltyBoxes: ***count subpopulations of network prefixes in various flap penalty states at any given moment.***

Objective: direct quantitative comparison of routing instabilities within their operational contexts.

Note: the PenaltyBox metric does not depend upon whether flap-dampening is actually deployed along the ASPATHs. It is just a useful, well-understood metric for measuring routing stability.

Guidelines for Fast Metric Design

1. Distinct peer routers have different views on global routing state, and send different temporal patterns of UPDATEs.

2. Must factor in peer consensus and temporal correlations across UPDATE streams.

3. Must factor out temporal jitter in outage start / end across peers.

Example: differences among peers – UPDATE patterns

row = peer

green mark:
announce

red mark:
withdraw

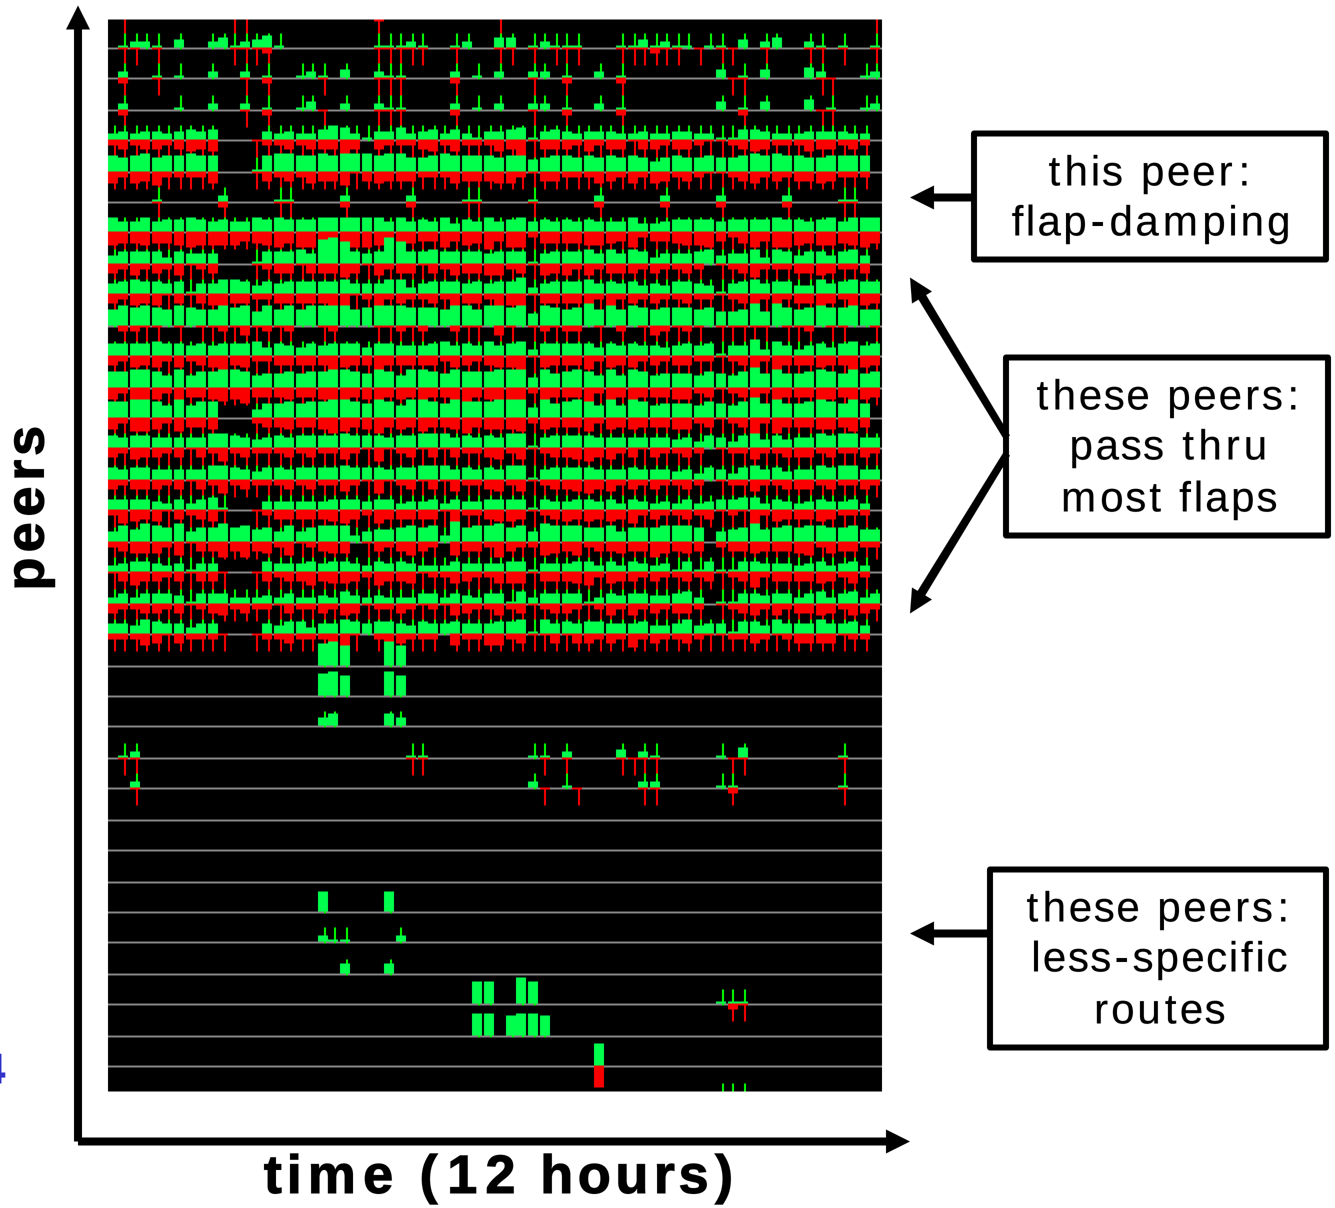

1 Feb 2004,
routes to one /24

Mechanism for such differences: distinct policies, paths

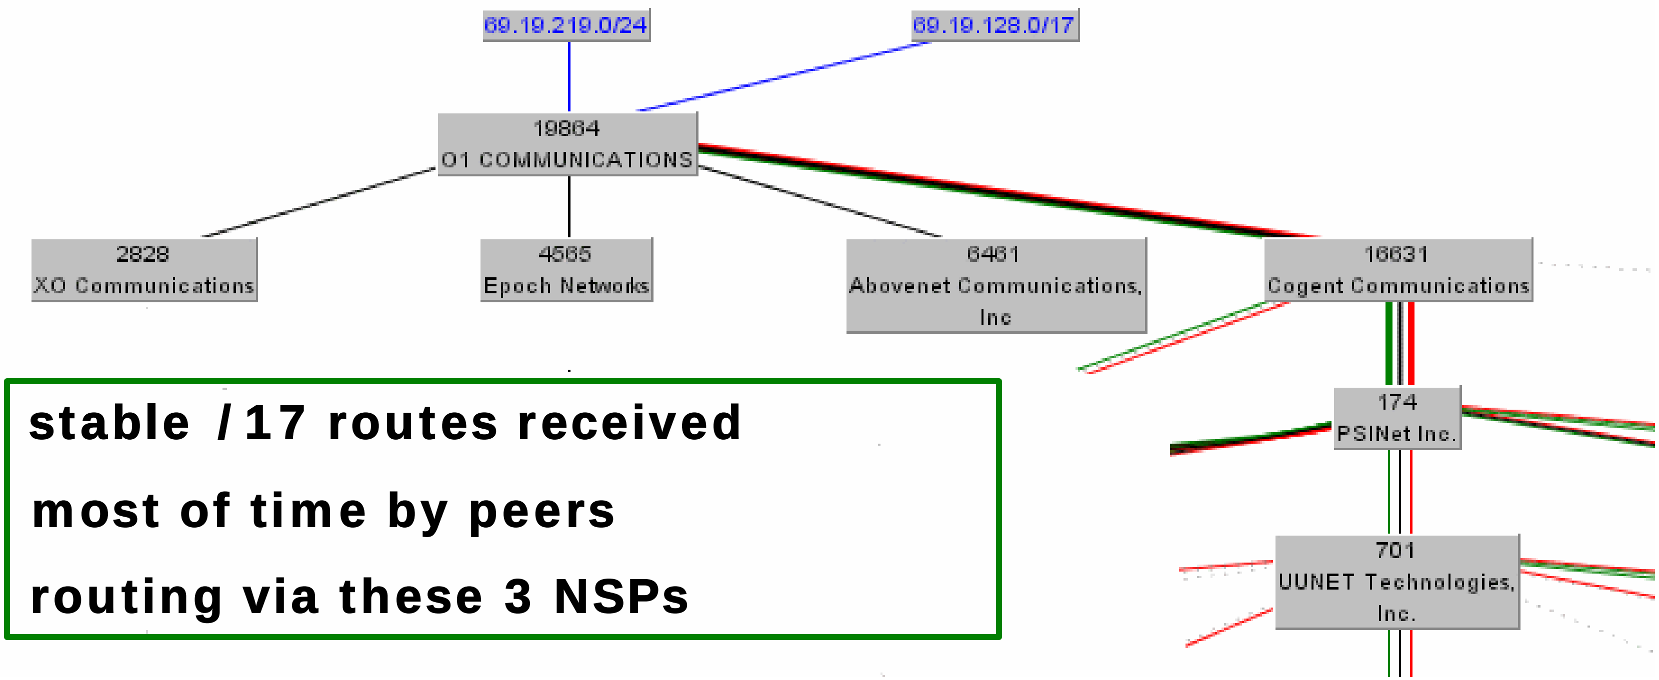

stable /17 routes received most of time by peers routing via these 3 NSPs

unstable /24 routes received by peers routing via these NSPs

flapping problem since 12/17/03

Example: differences among peers – flap damping outages

3 prefixes originated by one single-homed AS

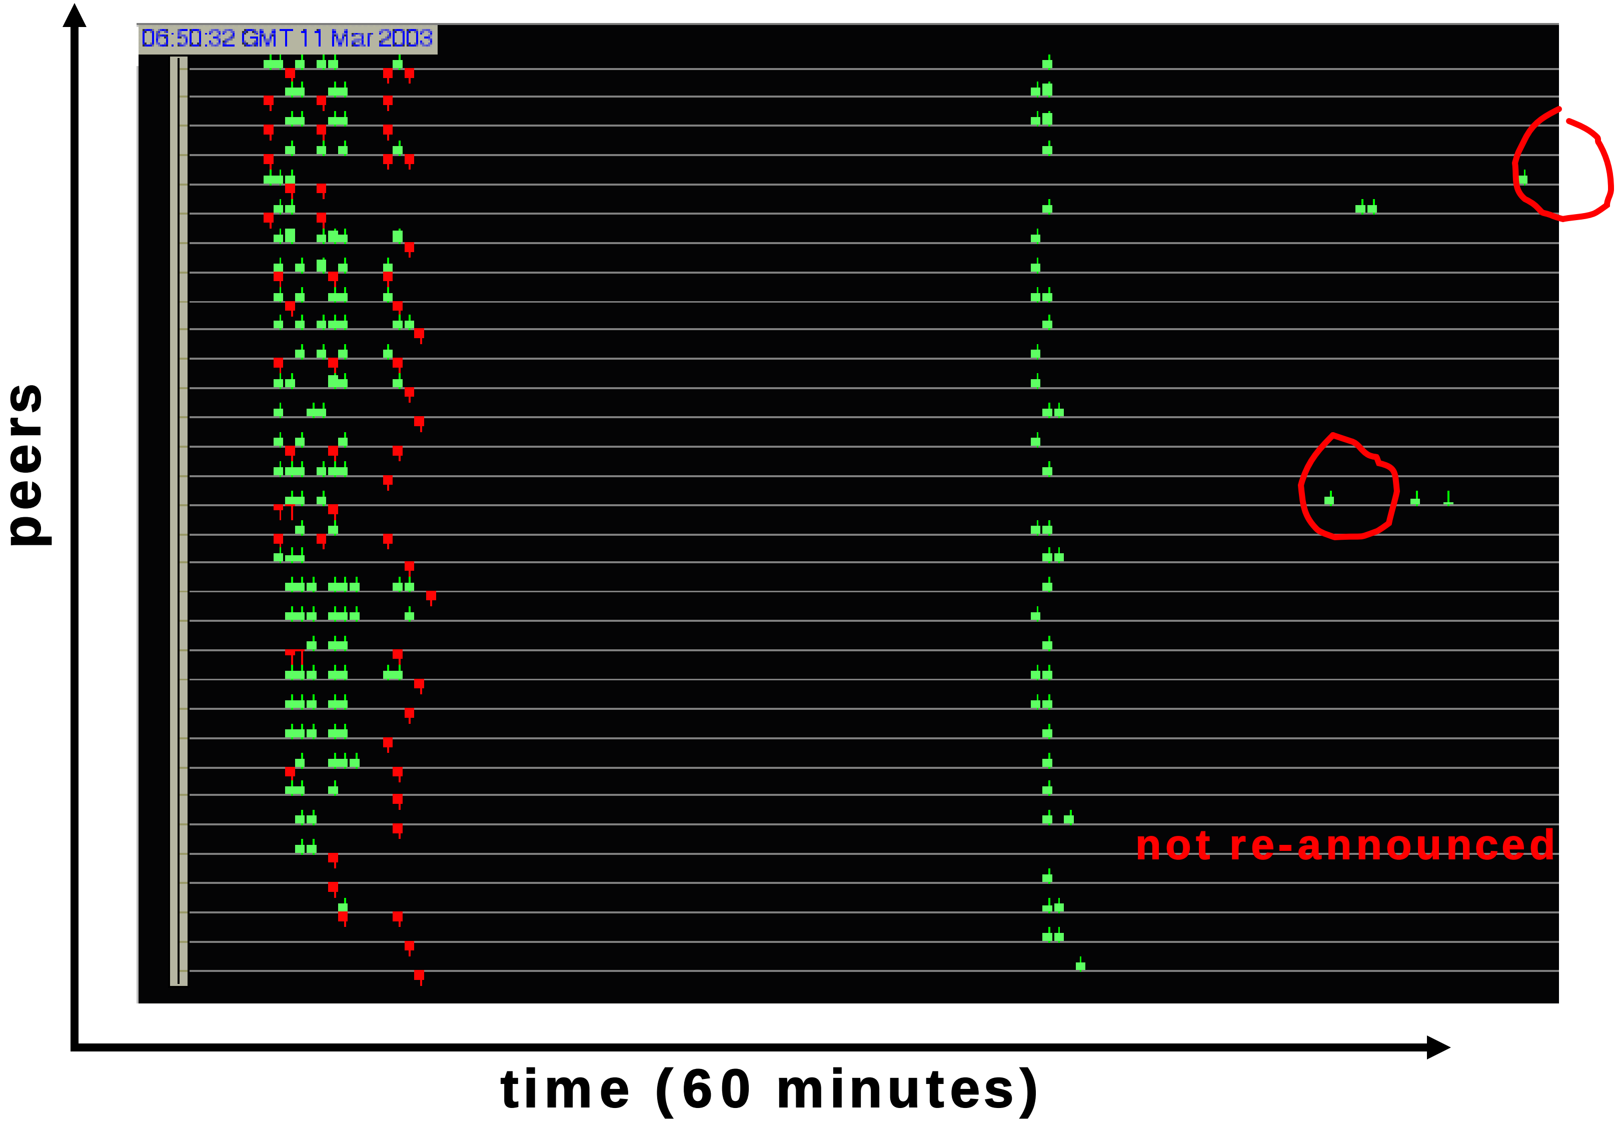

Guidelines for Fast Metric Design(2)

1. Discount outages reported by few peers; enhance outages reported by many peers.

2. Provide standard “knobs” for tuning the metrics; use different settings to filter for different kinds of phenomena, reproducibly.

3. Correlate further across subpopulations of prefixes common to a region, country, or set of origin ASNs; incorporate sensible strategies for averaging and normalization.

First Metric: NetsOut

NetsOut(T,P,D,MIN,MAX): the number of globally routable prefixes at time T that are suffering an outage.

A prefix outage event begins when at least P peers have seen withdrawals for some prefix, separated by no more than D seconds between successive withdrawals, with no new advertisements interspersed. It ends when fewer than P peers continue to believe the prefix is withdrawn.

Disregard outages that are shorter than MIN seconds or longer than MAX seconds in duration.

Reasonable: P={3 -10}, D={60s - 120s}, MIN=120s, MAX=7d

NetsOut sensitivity to parameter choices

1. peer withdrawal jitter D

(30 to 120 sec)

sensitivity = SMALL

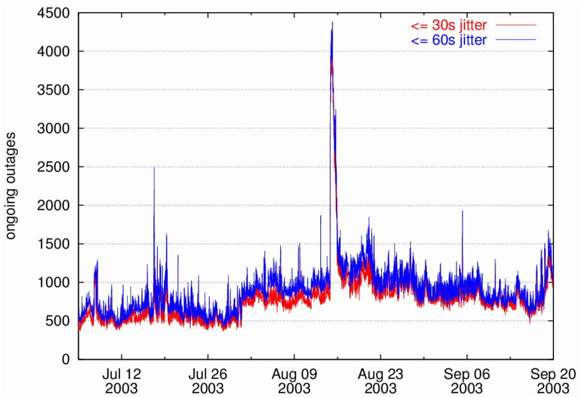

2. correlative peer agreement count P

(over 3 to over 30 peers)

sensitivity = MODERATE

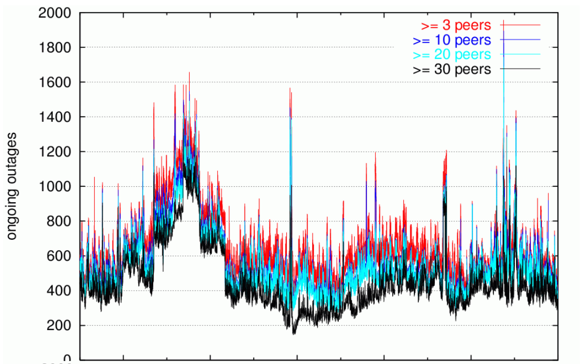

3. maximum outage duration MAX

(0.5 to 6 days)

sensitivity = NOTABLE

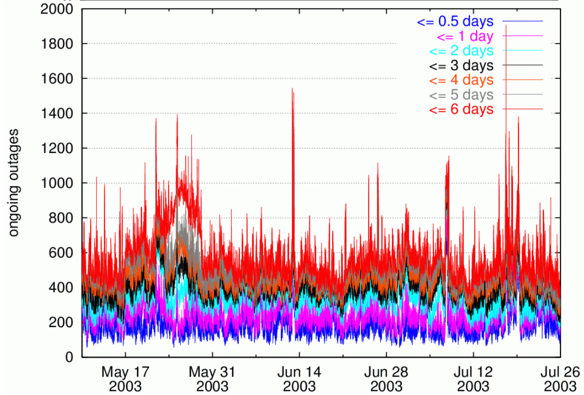

reason: prefix churn

Long-term view on concurrently ongoing transient prefix outages (NetsOut) 3Q 2003

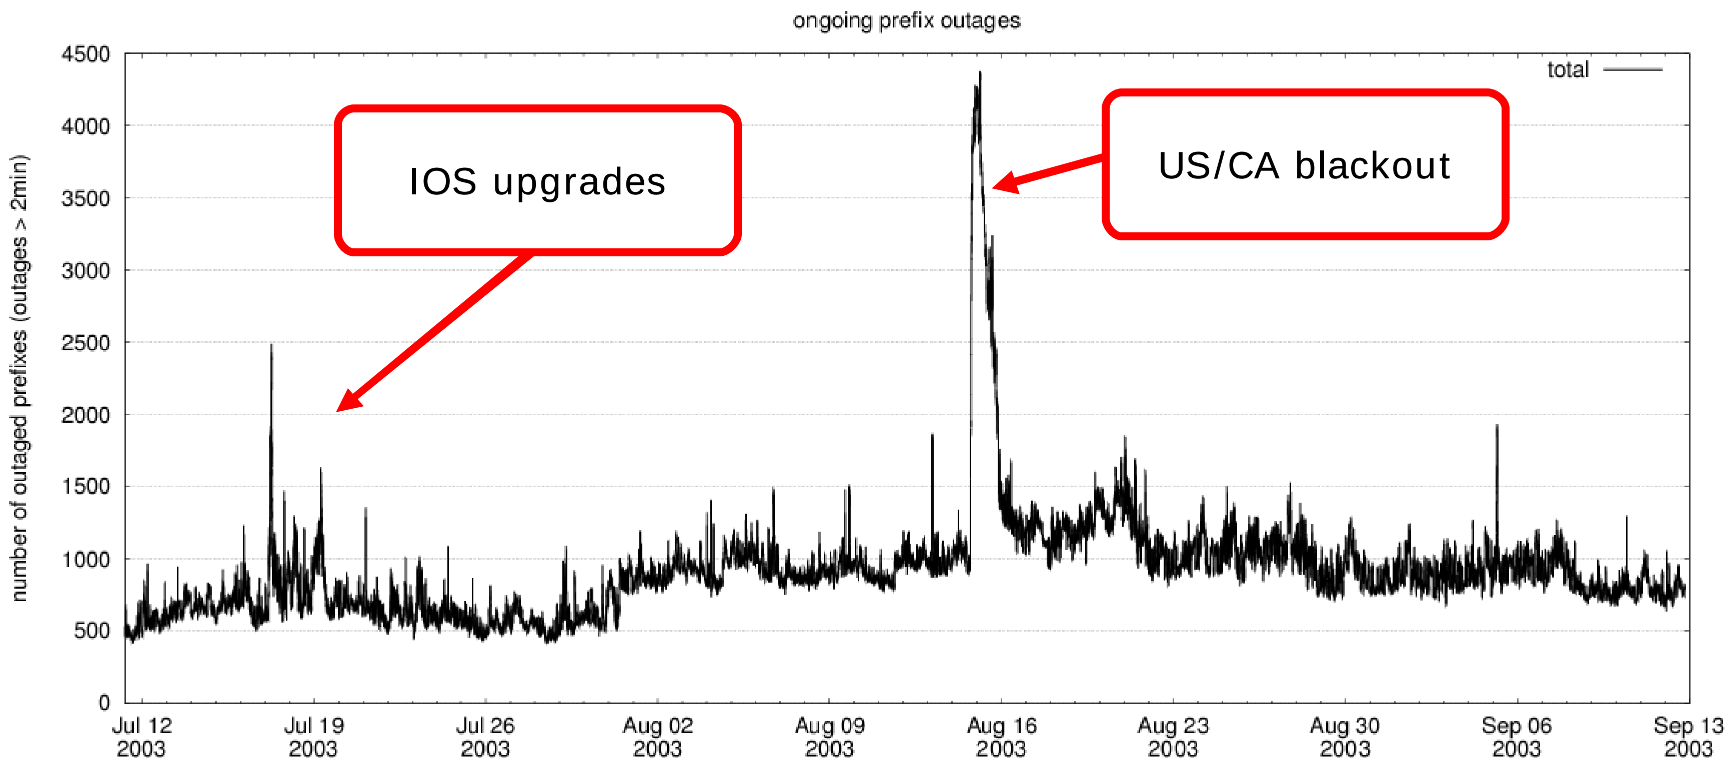

P=3, D=60, MIN=120s, MAX=7d

Distribution of outage durations: 3Q 2003

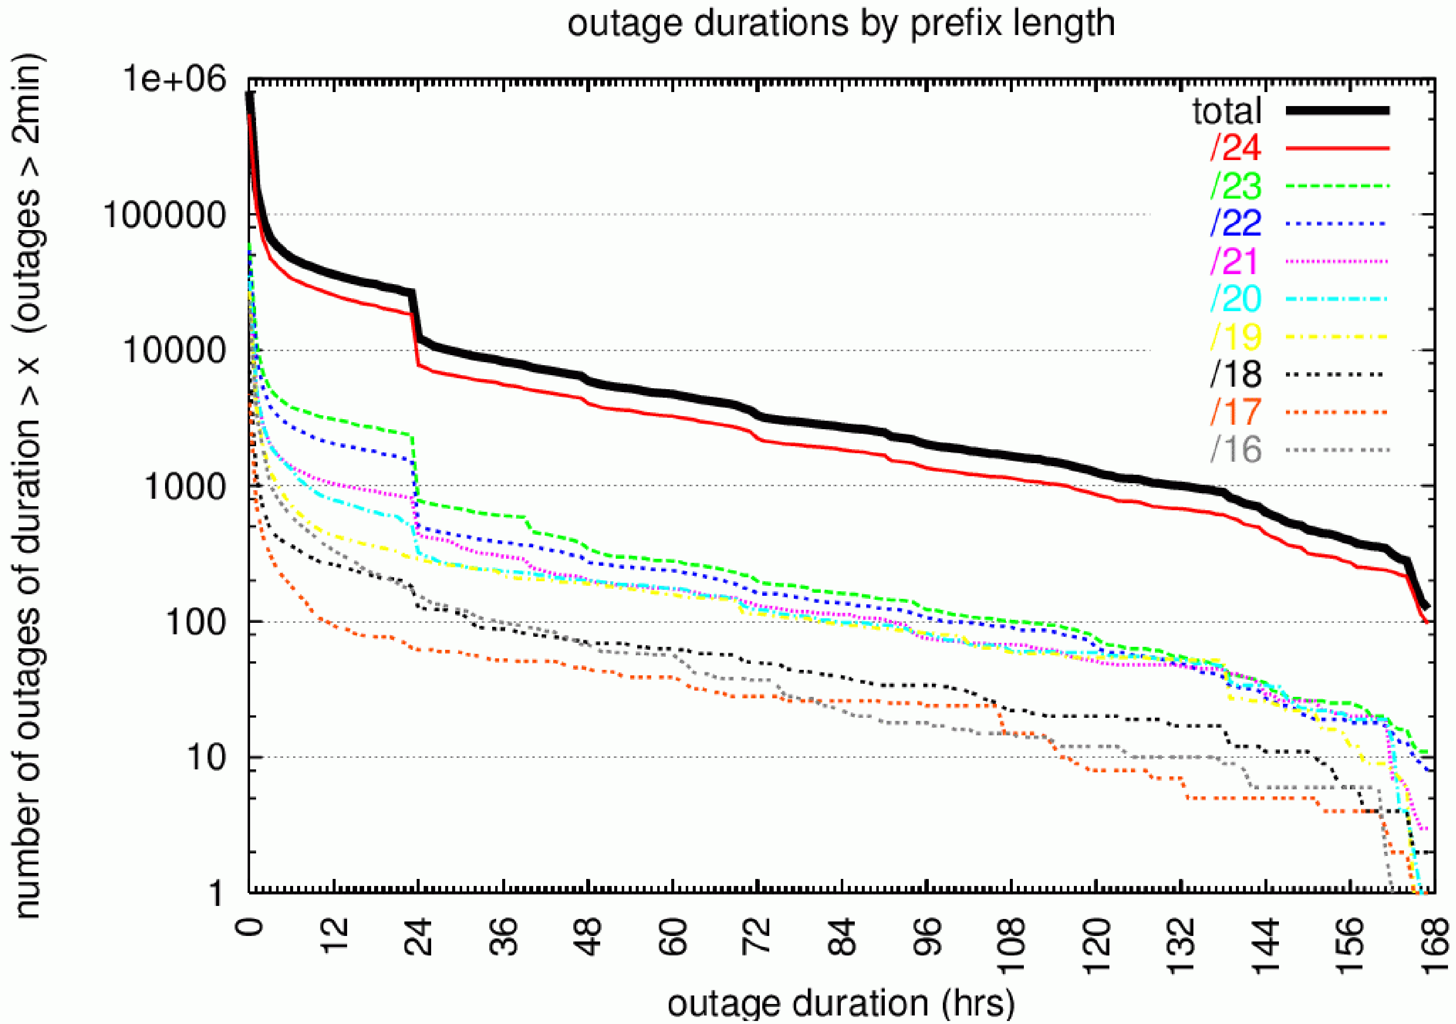

Second Metric: PenaltyBoxes

PenaltyBox(T,K,H,C): the number of globally routable prefixes at time T that have flap penalty K, using the classic flap dampening algorithm with half-life of H and ceiling of C.

PenaltyBox K across multiple BGP sessions at time T is the arithmetic mean of the number of prefixes with penalty K across all peer routers.

Apply to all prefixes in full BGP tables for global metric, or subsets (countries, single ASes) for local metrics.

Reasonable: H=600s, C=15, K={0,...,C}.

Ongoing research: sensitivity to choice of dampening algorithms (vendor C vs vendor J..., RIPE recommendations.)

flapping prefixes in PenaltyBoxes – 3Q 2003

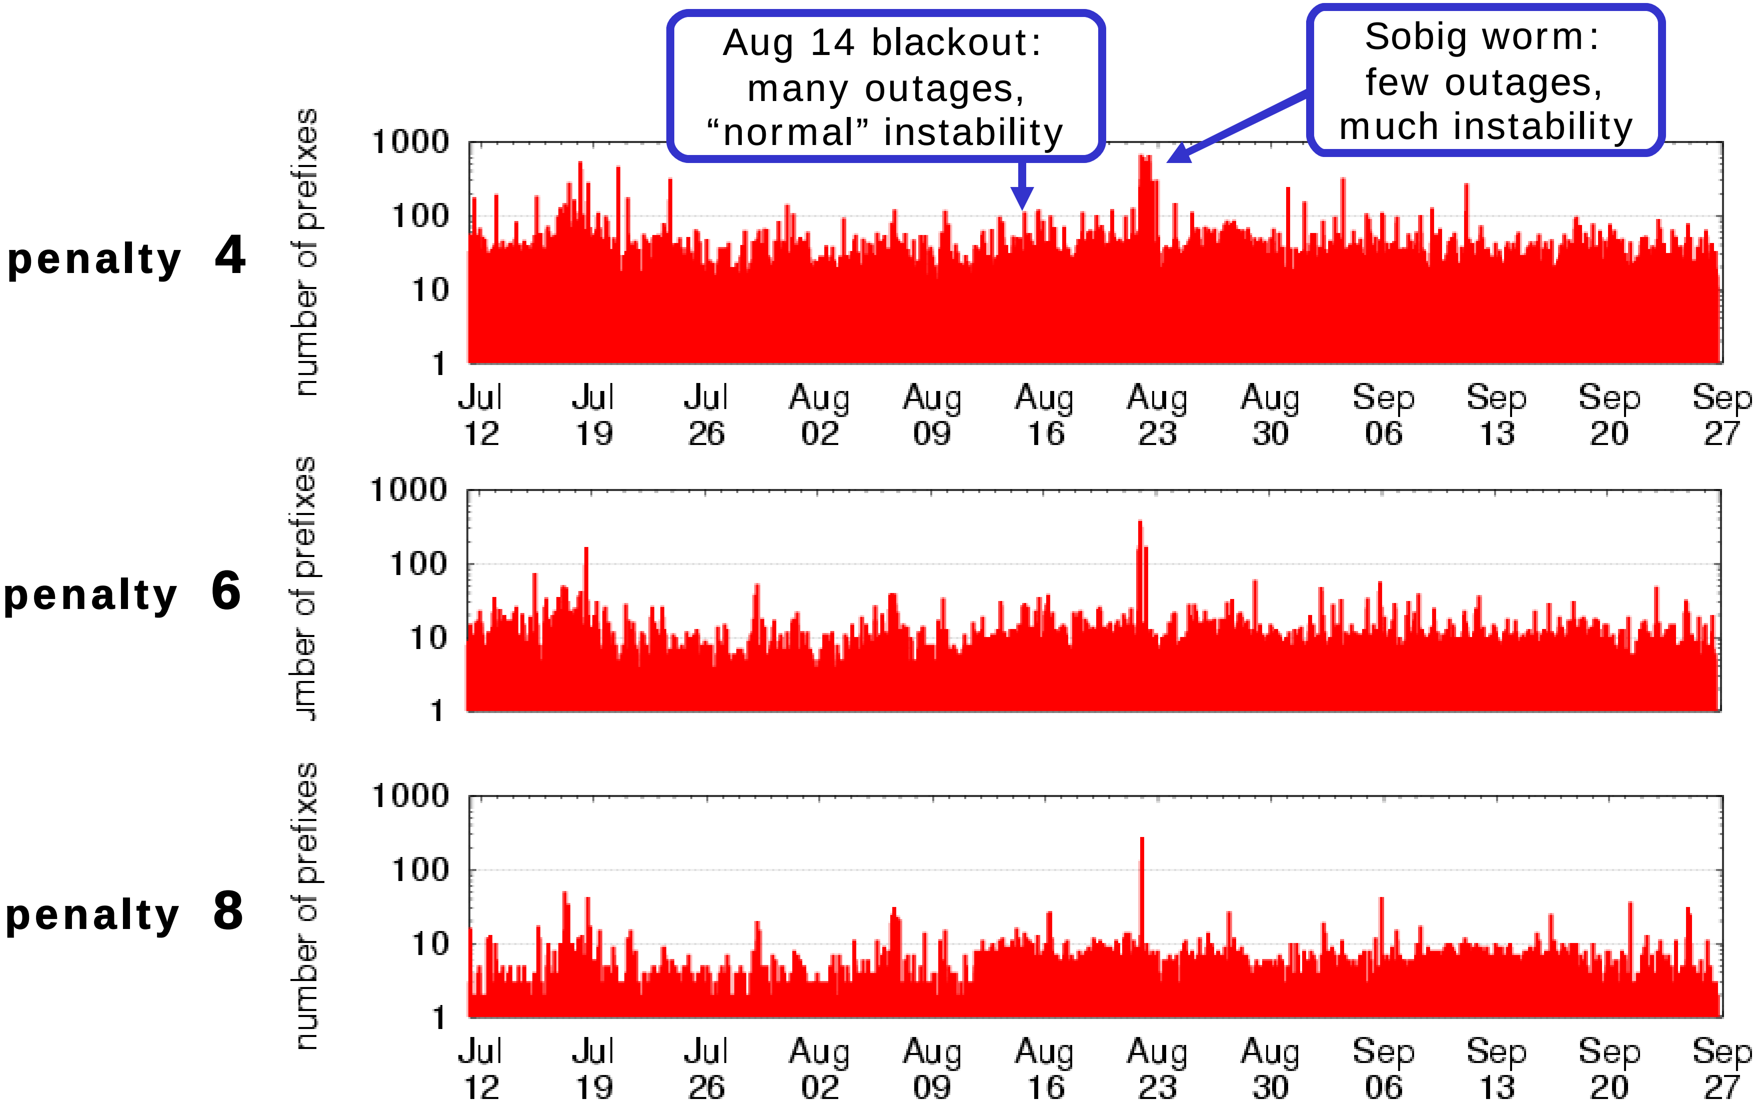

A 3d view on all PenaltyBoxes: long flapping AS

Prefixes originated by one NSP, Aug-Nov 2003

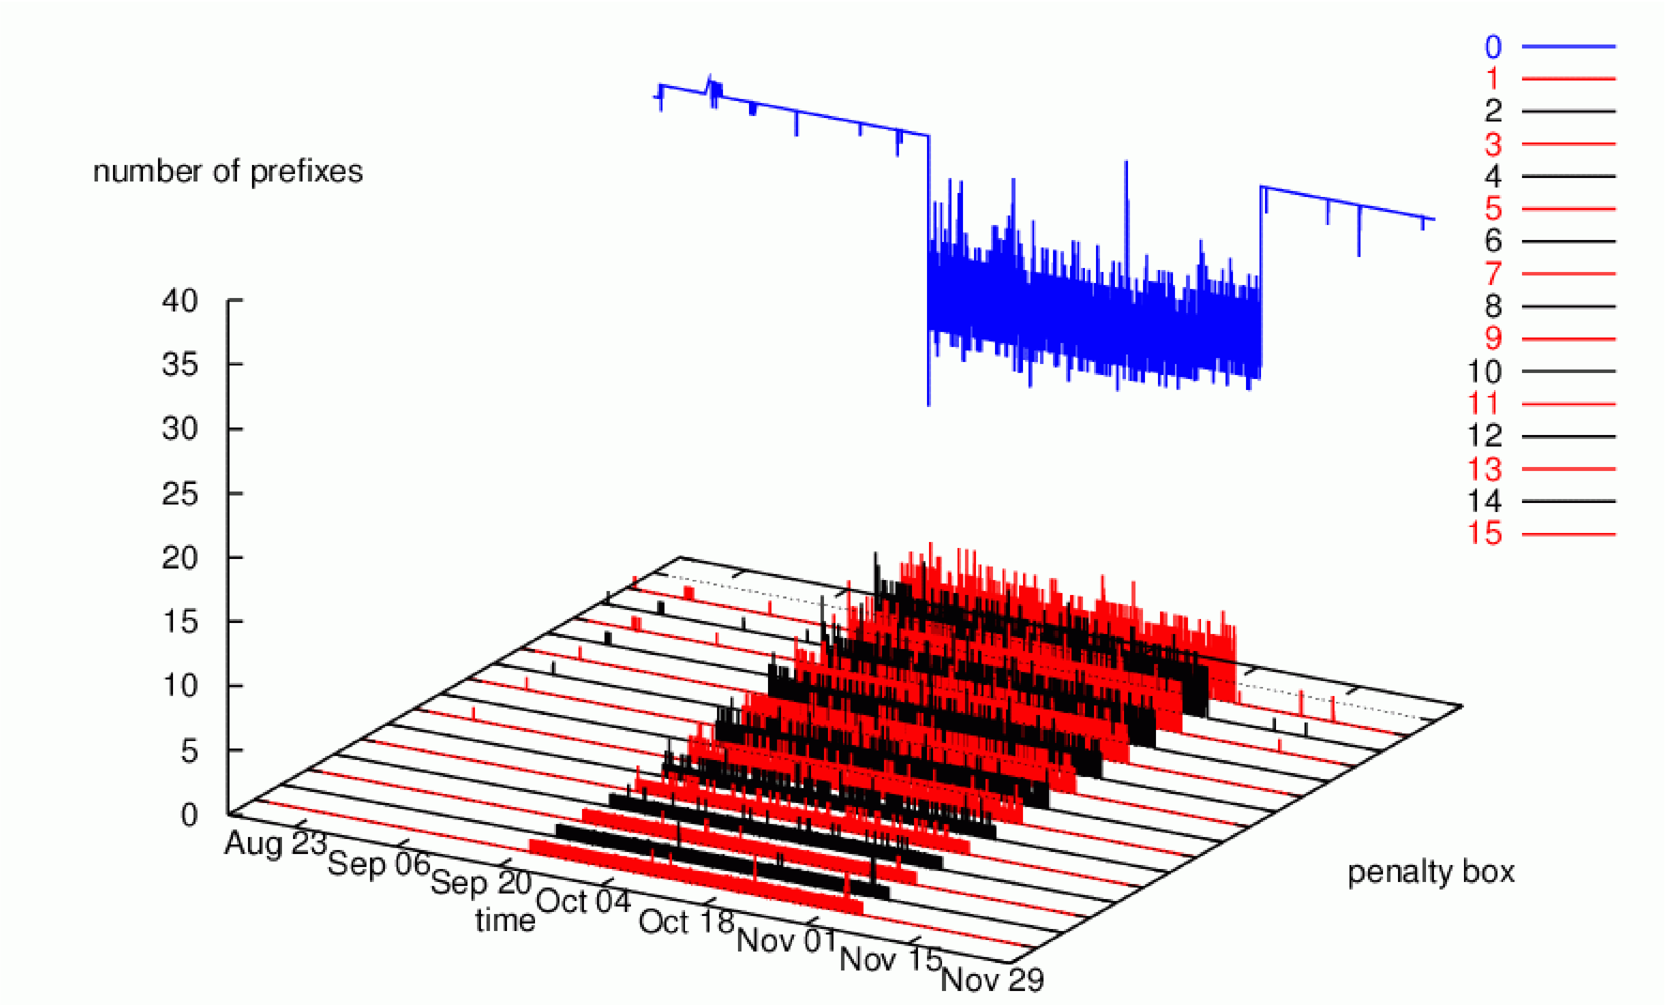

Why bother with all this?

Long-term measurements of the numbers of transiently withdrawn networks, and of prefixes with widely-seen flapping routes are an obvious thing to do. So…

…Are the measured instability levels acceptable?

…Why do well-respected ASes often flap routes for very long periods of time, undetected?

… Why does the operations community still lack an objective standard for comparing the impact on global routing of more/less serious events (SCO DoS vs Slammer)?

… How large do “excursions from the norm” normally get for these metrics, when the Net is perturbed in these ways, and

… *How large can such excursions get?*

Conclusions:

Fast Routing Instability Metrics

- Simple, intuitive metrics can capture and quantify even very high-dimensional, diverse routing behaviors.
- Metric sensitivity and robustness depend on selection of good parameters, and on integrating multiple diverse data sources to get the big picture.
- Not all threats impact the routing infrastructure; good metrics help us automatically distinguish “layer 8” effects (SCO-type viruses) from real infrastructure threats (warhol worms and router DoS).

Conclusions: Fast Routing Metrics (2)

- These are metrics you can compute for yourself, over your own data. BGP UPDATEs are freely available from RIPE, RouteViews, or your friendly local BGP-speaker.
- We gladly accept BGP feeds to increase the diversity of our sampling space, and the statistics we report to the community.

  http://www.renesys.com/peering

- Thanks!

Additional slides

NetsOut sensitivity to peer withdrawal jitter D

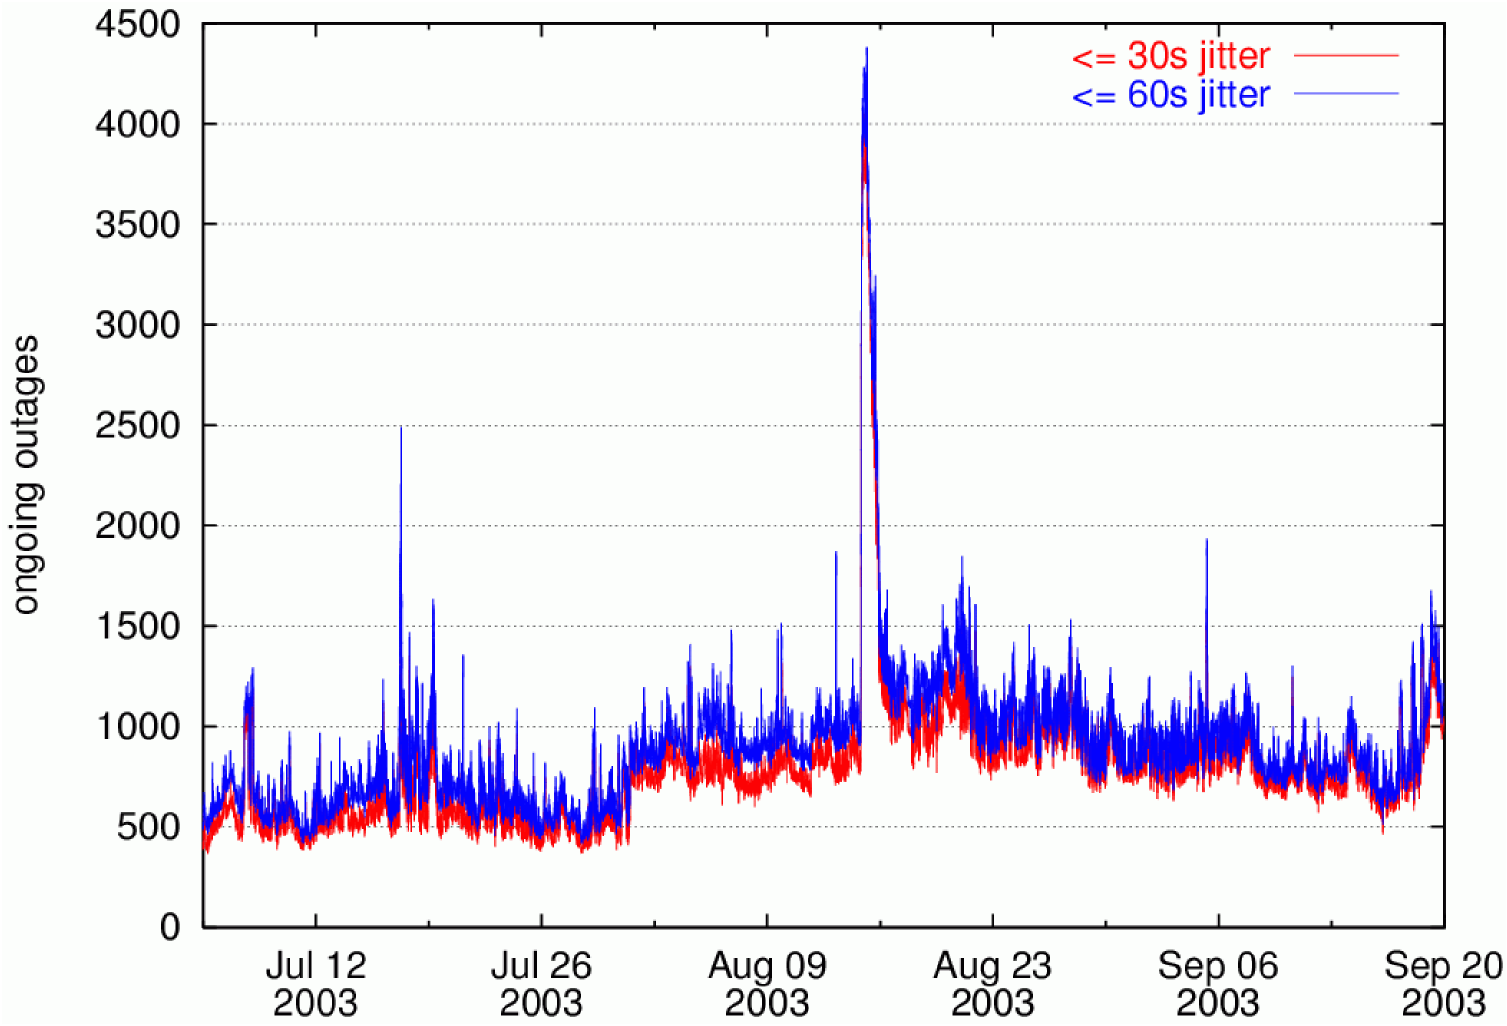

P=3, MIN=120s, MAX=7d

NetsOut sensitivity to correlative peer agreement count P

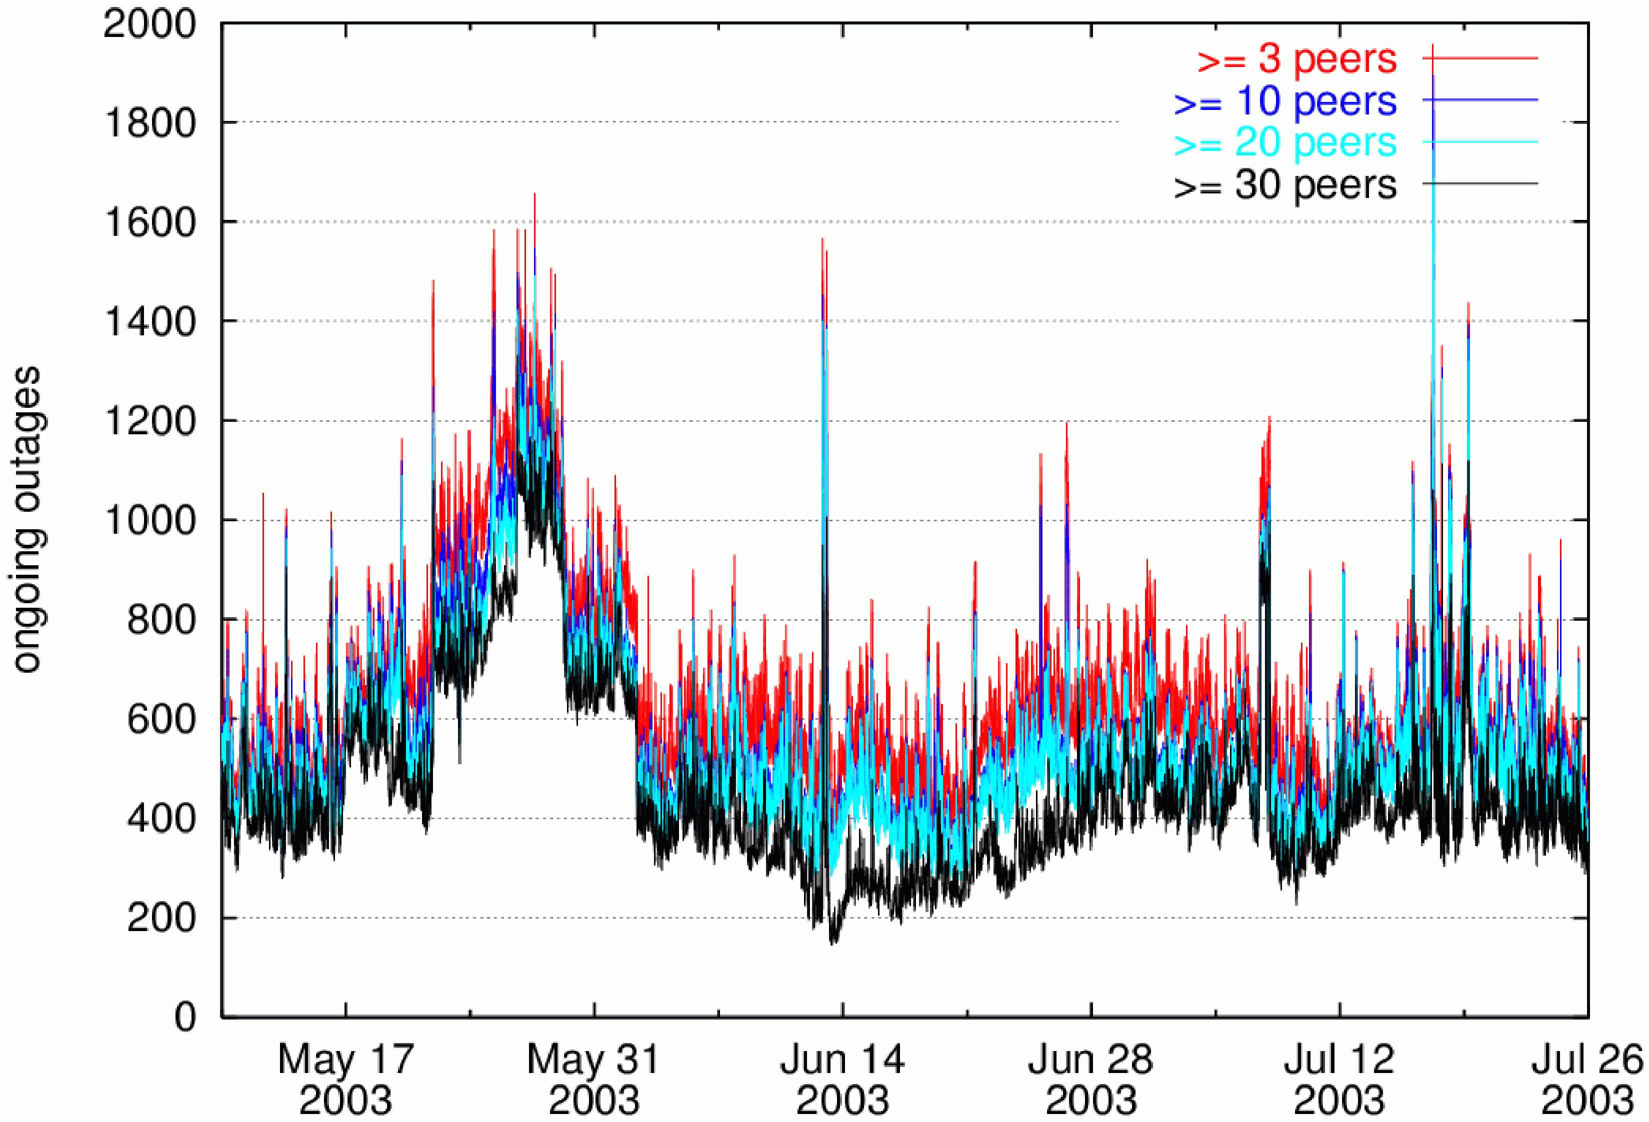

D=30, MIN=120s, MAX=7d

NetsOut sensitivity to max outage length MAX

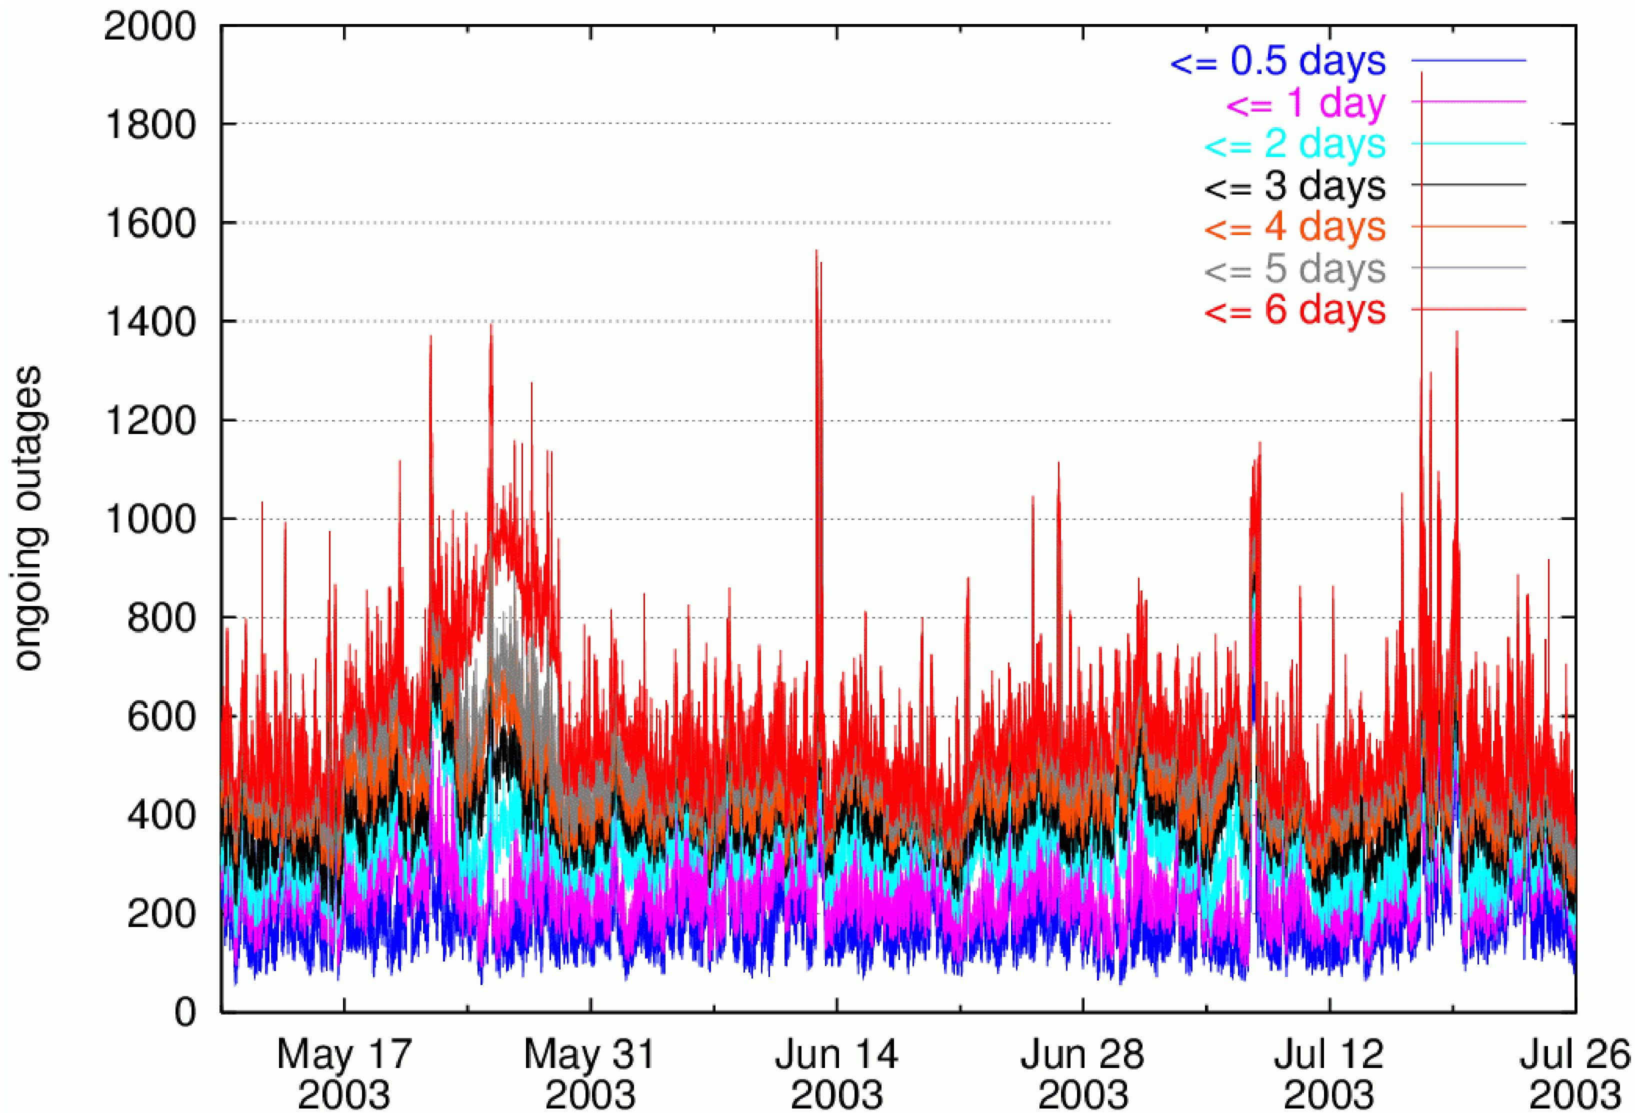

P=3, D=30, MIN=120s