Some thoughts on CoS and Backbone Networks

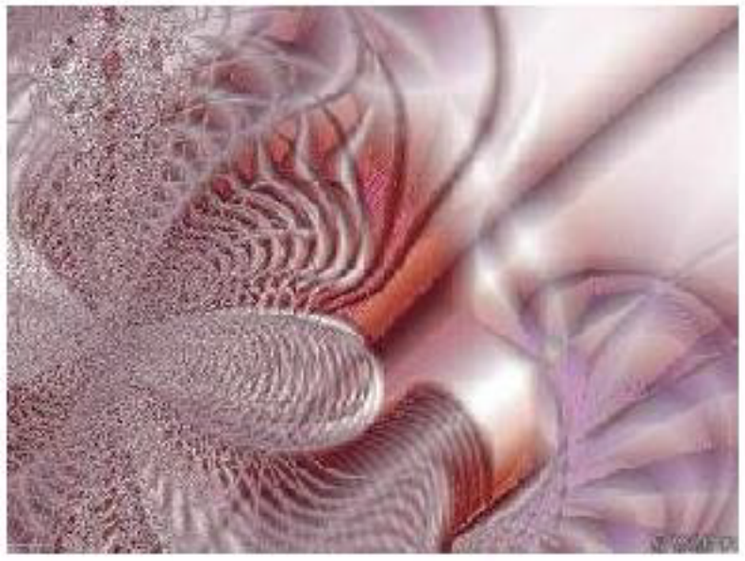

David Meyer
ieprep, 55th IETF
Atlanta, GA

Agenda

- What Sprint is doing
  - What it is
  - What it isn't
- Is QoS is required in an internet core?
- http://www.maoz.com/IETF55/ieprep

What Sprint is Doing

- Idea: "Edge QoS" (aka eQoS)
- Start with a congestion free core (no queues)
- Do queuing (customer configured) on the edge box
  - Only on egress from our network
  - Using ACLs to avoid ToS bits (except for IPSec VPN)
  - IPSec VPN -- case TOS bits
    - If those bits encrypted --> per CoS tunnels
    - Need ToS bit export (to IPsec header) to use single tunnel
- Note "problems"
  - Ingress v. egress
  - ACL scaling
  - DoS

What Sprint is Isn't Doing

Respecting any bits other than destination address in the core

Why you don't need QoS in the core

- Start with the result that aggregated traffic in the core (>= OC48) is "uncorrelated"
  - "Impact of Aggregation on the Scaling of Internet Backbone Traffic", Zhi-Li Zhang, et. al, Sprint ATL Technical Report TR02-020157 (http://www.sprintlabs.com/ipgroup.htm).
  - Many others...e.g., Thomas Telekamp's talk(s) from the last NANOG (http://www.nanog.org/mtg-0210)

- So you can actually build (provision) a network to avoid queuing
  - Oh, and BTW, if you do this, you get close-to-zero packet loss, close-to-speed-of-light delays, and small jitter
  - With no additional mechanism!

So what would QoS mean anyway?

- Caveat: The following graphs are for a system having a Poisson arrival process and exponentially distributed packet sizes (so I could use gnuplot :-))
  - And...while this isn't a perfect model of Internet traffic, it is a reasonable approximation and makes the associated math much easier (M/M/1 model)
    - http://ipmon.sprintlabs.com/paccess/tr/public/TR02-ATL020312.pdf has supporting empirical results
- So... they are meant to show the effect of increasing bandwidth on queuing delay
  - And NOT to be a statement about actual queuing delay
- As such, the following graphs do a good job of showing the general behavior

Queuing Delay vs. Utilization, T3

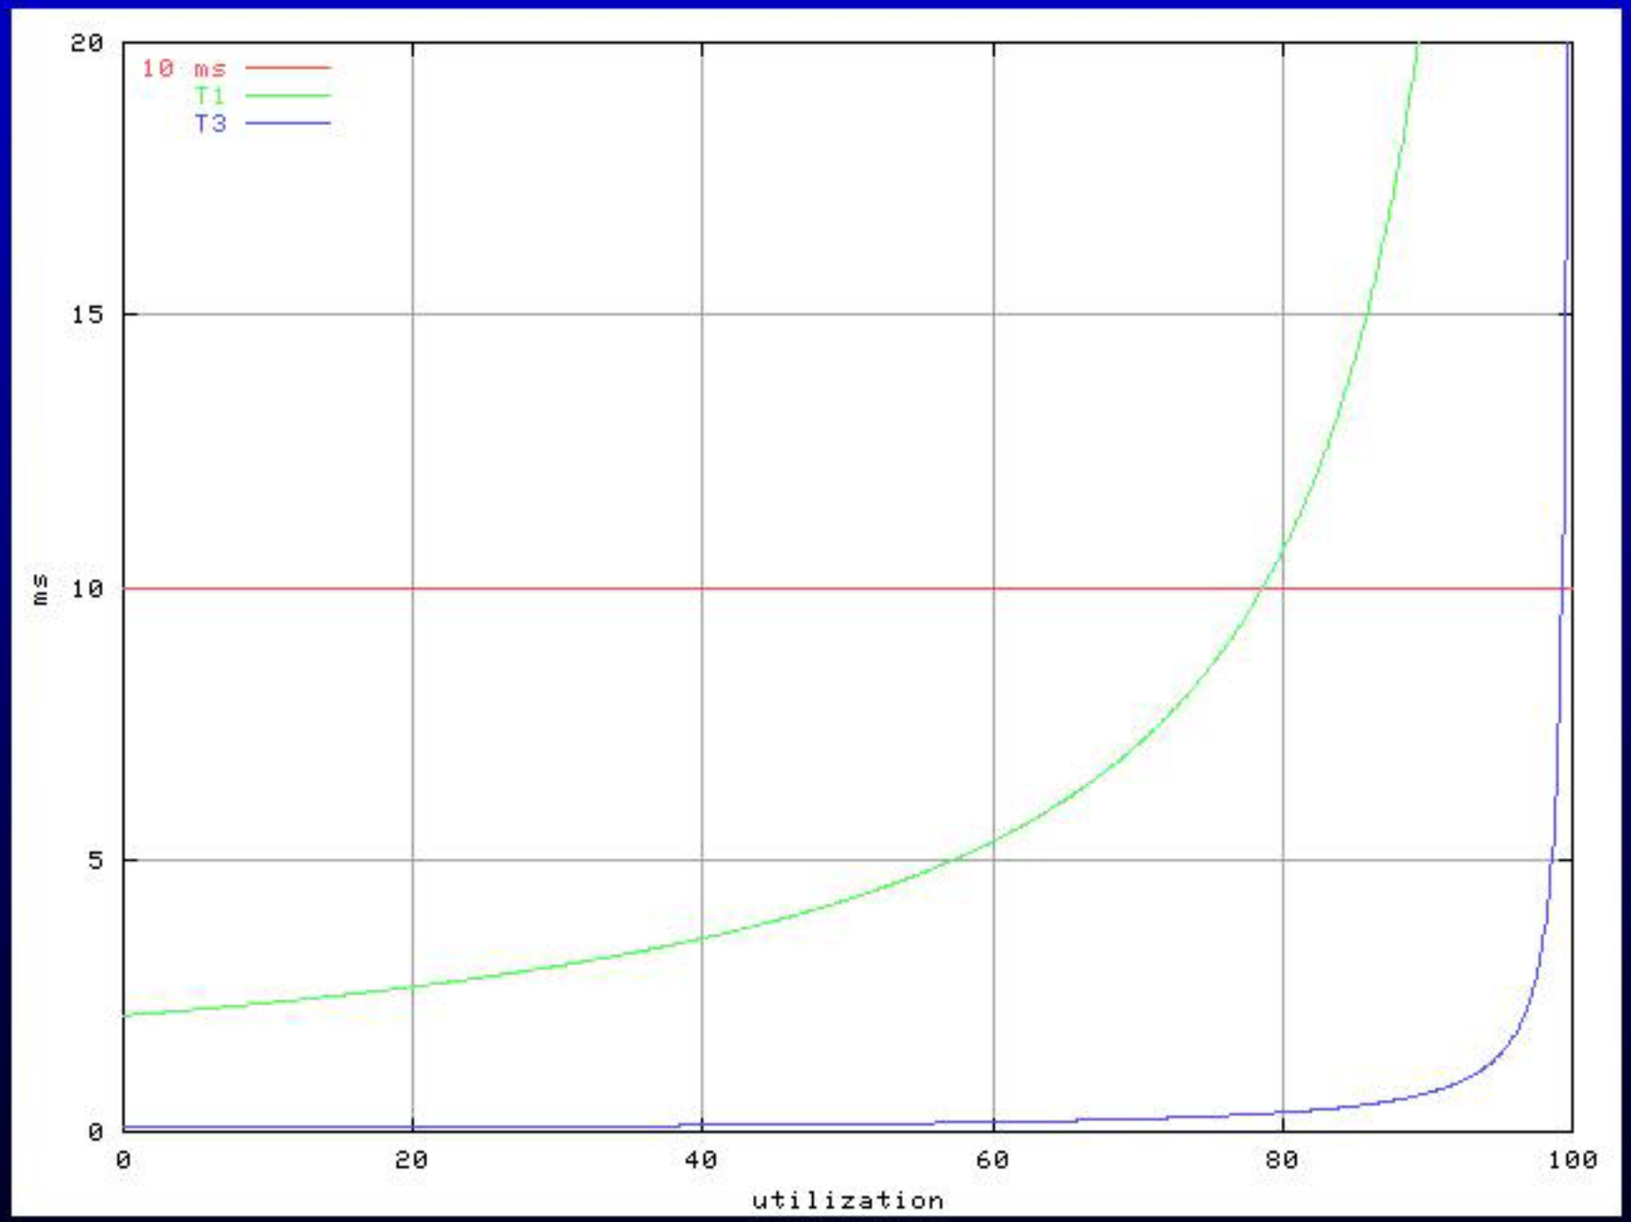

Queuing Delay vs. Utilization, OC12

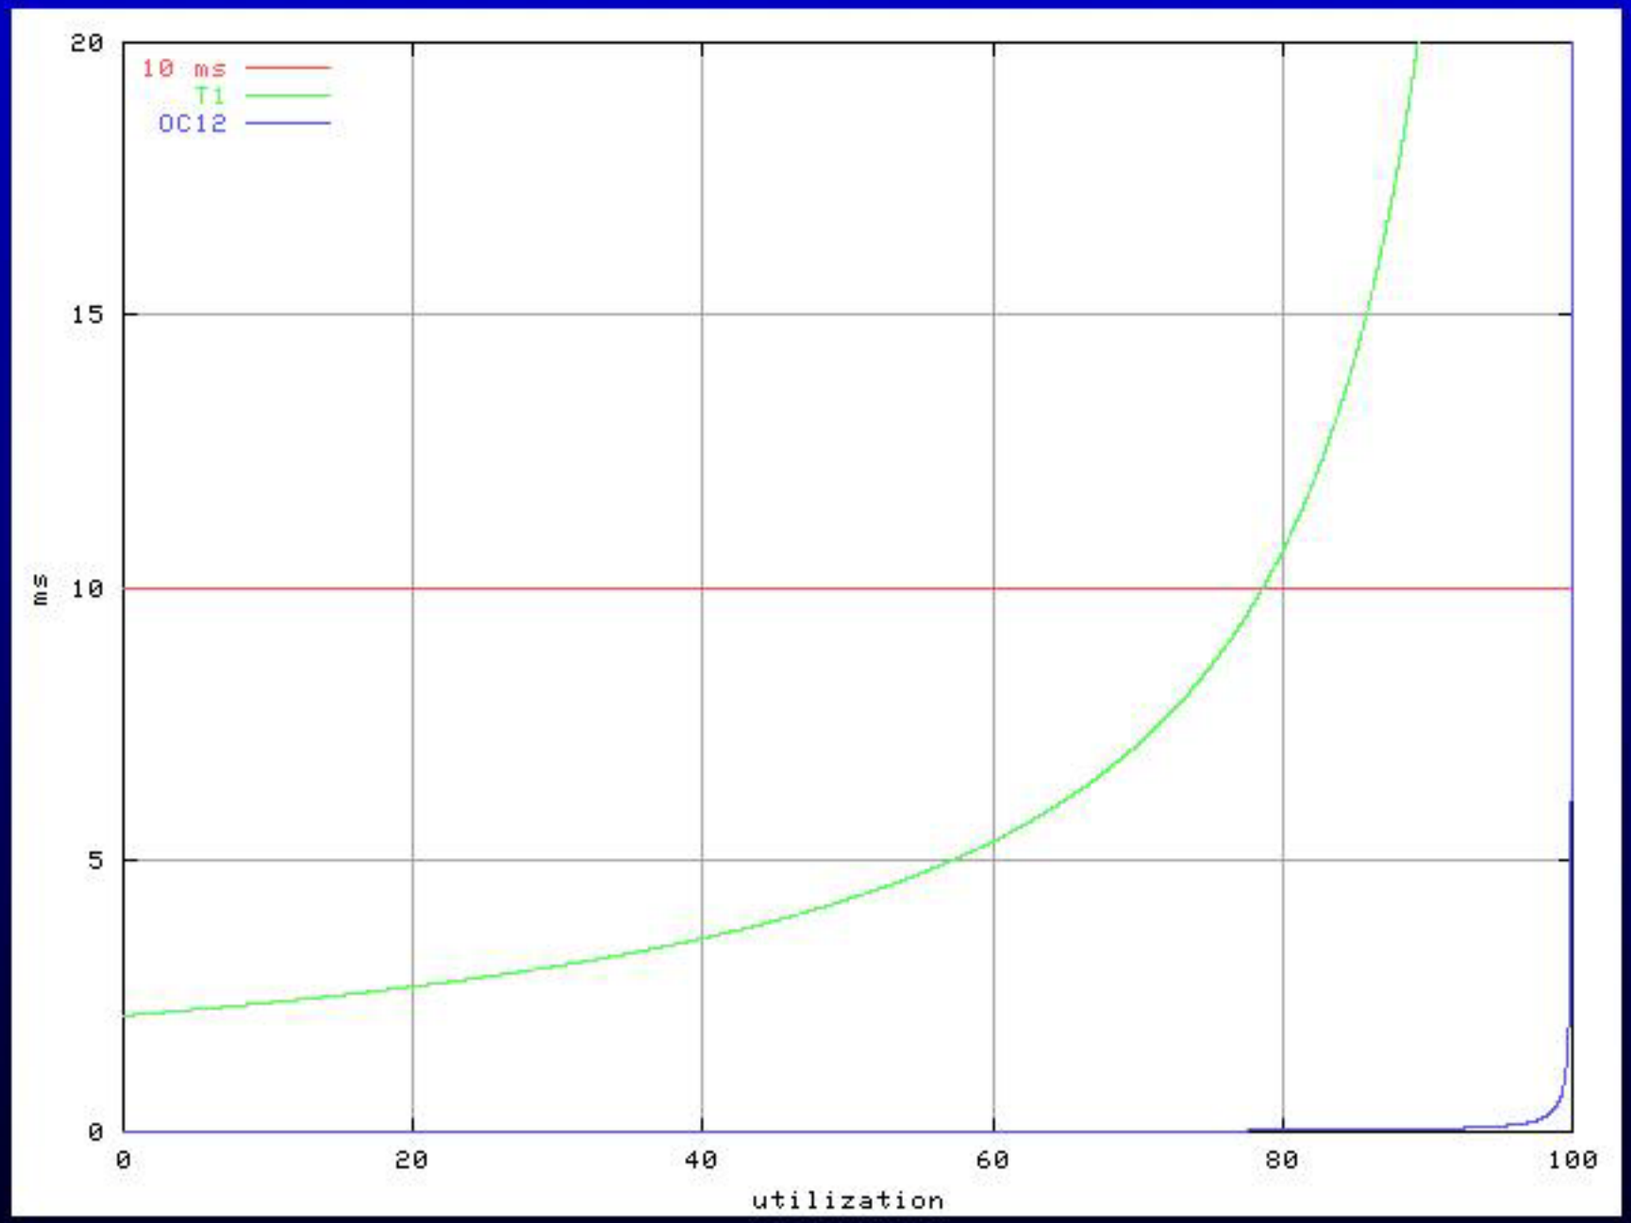

Queuing Delay vs. Utilization, OC48

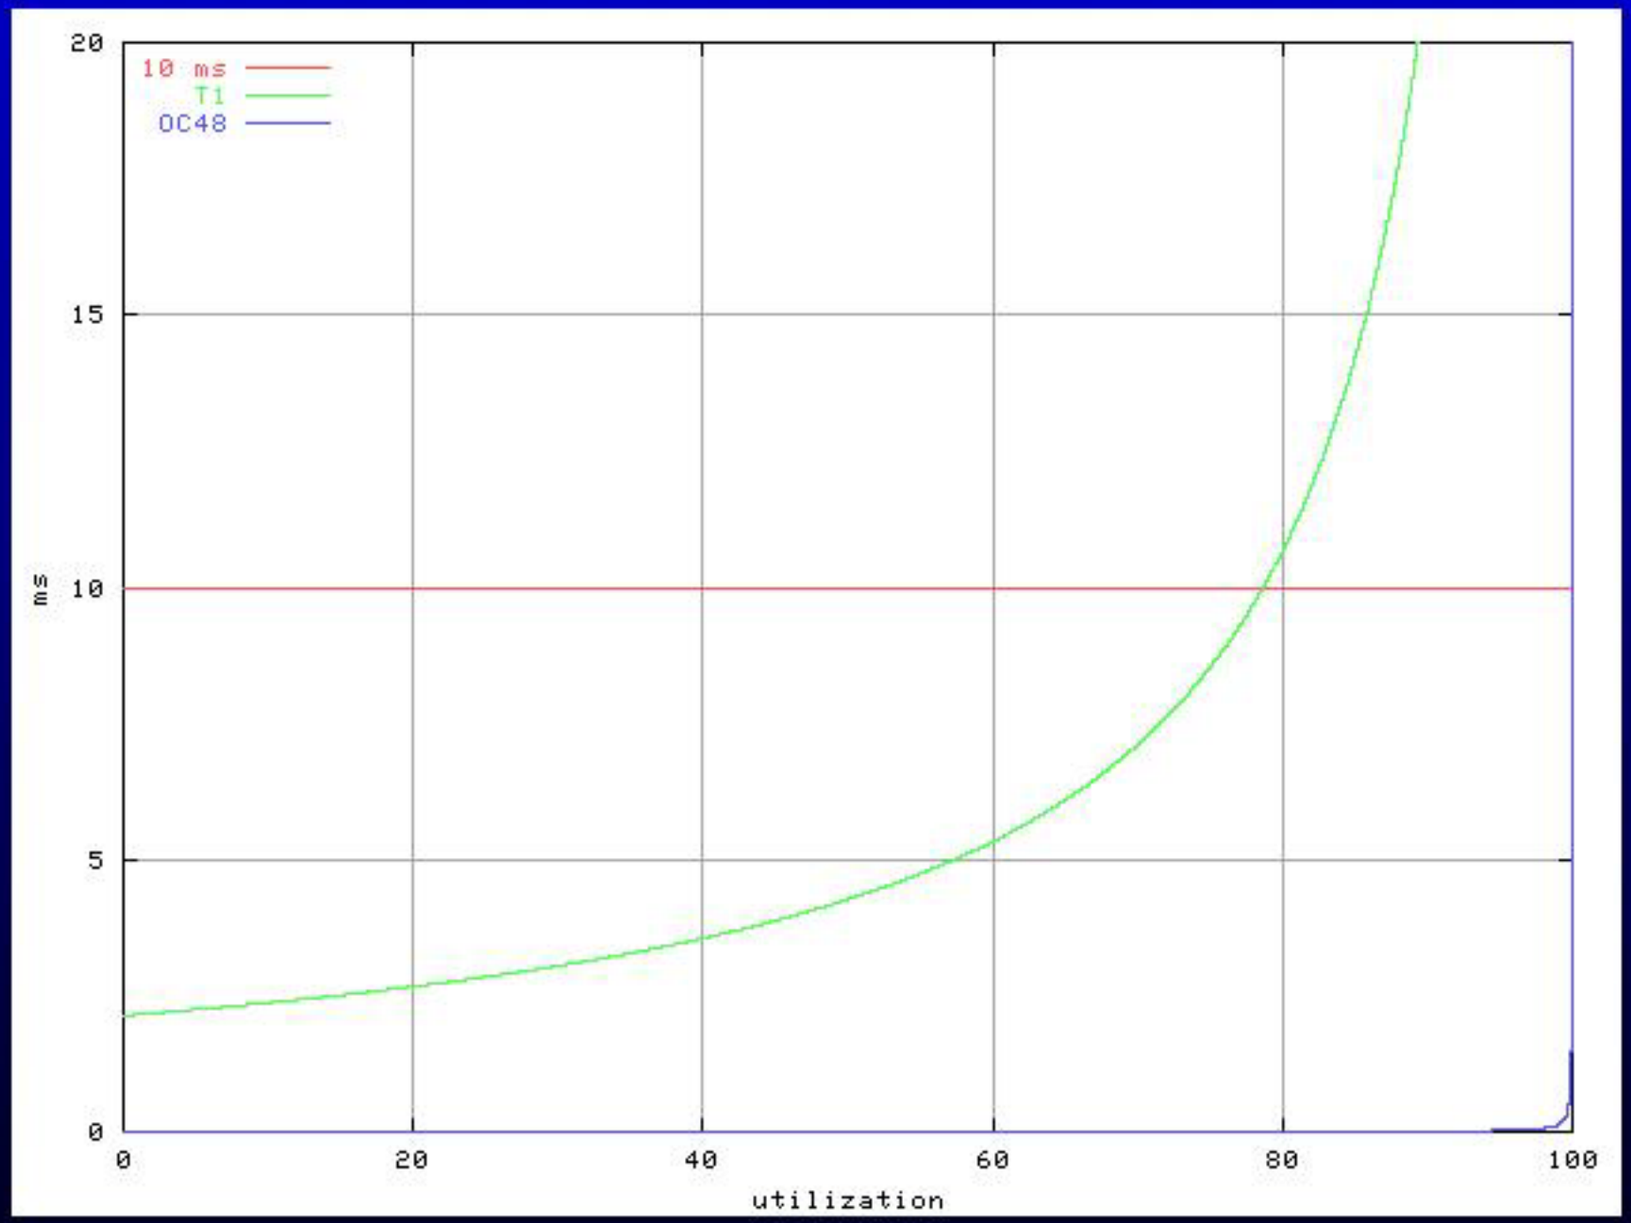

E-model ratings/MOS Scores

- We studied the ITU-T E-model (G.107, a subjective measure of call quaility based on measurable parameters (delay, loss, jitter))
  - Range: 100-0, with 100 best, 0 worst
  - PSTN: 100-70
  - Cellular: 90-60

- Briefly, methodology:
  - Calls are placed at random times (Poisson arrival process), and we emulate business call duration (~ 3.5 minutes)

- In the following graph...
  - The blue line is the average call rating
    - 1 call/minute for 24 hours
  - The red line is the 99.9% confidence interval

E-model ratings/MOS Scores

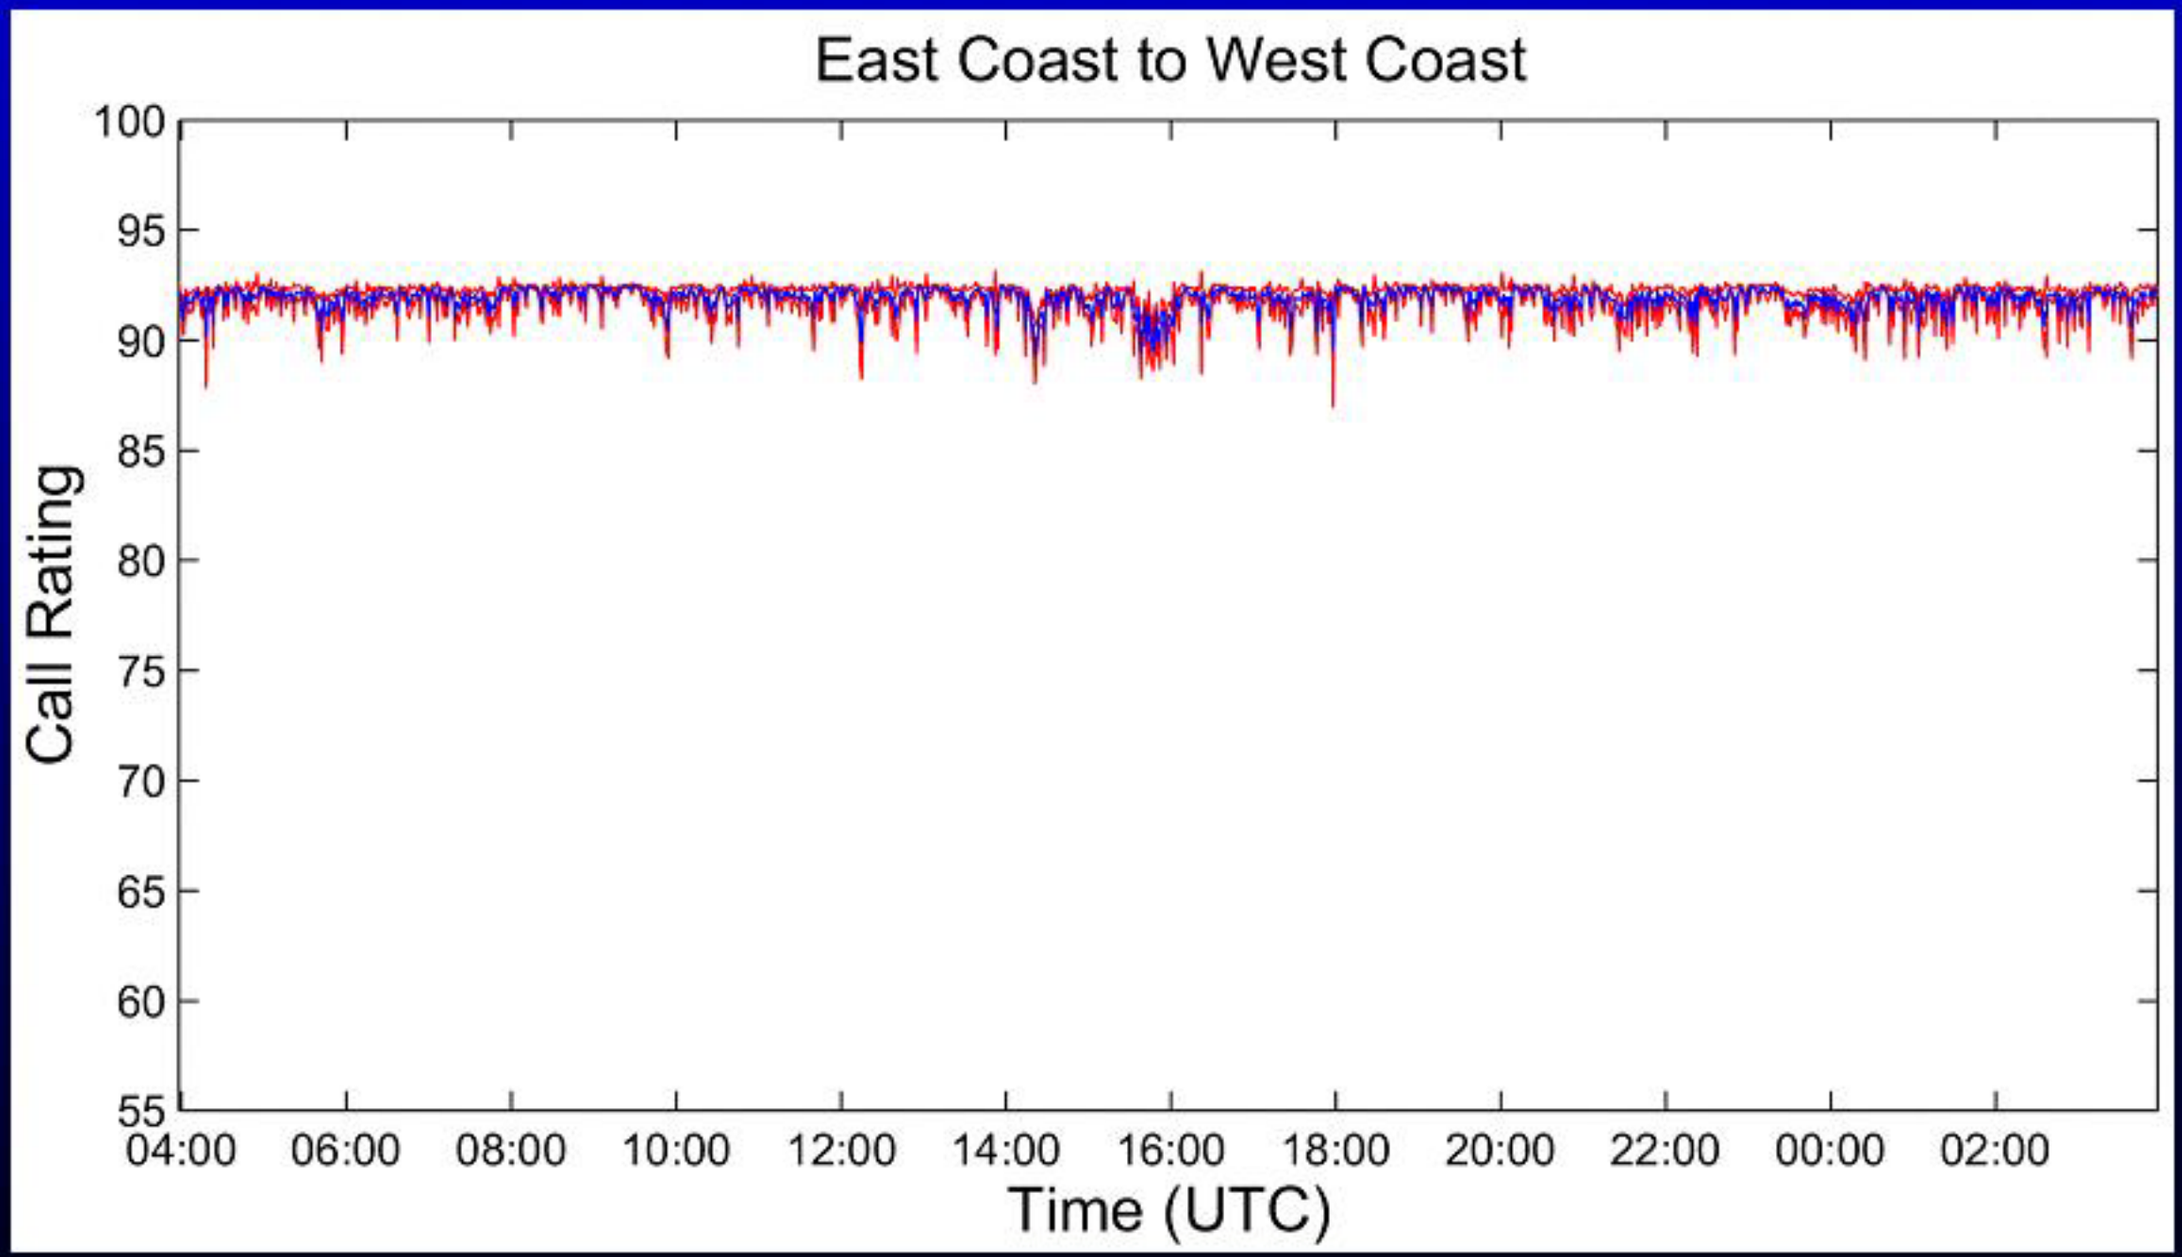

Questions/Comments?