RTTs for anycast root nameservers

Nevil Brownlee, CAIDA / The University of Auckland

20 July 2005

---

1. Overview

- ***Root RTT strip charts***
  - **Change in G root, 2 Jul 05**
- ***Root RTT sequence plots***

- ***Can root RTT changes indicate responses from different anycast instances?***
  - **We see bimodal behaviour at Auckland for K and M**
  - **We don't see it at all at Boulder**
  - **We did see it at Auckland back in 2002 —**
    we believed it indicated load balancing on root paths

- ***Unusual RTT 'steps' at Colorado, 9 Jul 05***

---

1. CAIDA's root/gTLD plots

- ***RTT data recorded as described in "Response time distributions for global name servers," Brownlee & Ziedins, PAM 2001***
  - **First 100 RTTs in a 5-minute interval are recorded, without timestamps**
  - **When more than 100 RTTs are observed in a 5-minute interval,**
    we bin them so as to collect an RTT distribution

- ***We produce two types of plot from our data:***
  - ***Strip charts*** **(on our root/gTLD web page) show the *median* RTT for each 5-minute interval**
  - ***Sequence plots*** **show individual RTTs in sequence, plotted at equal spacing within their 5-minute intervals**

---

1. Example root RTT median plot ('strip chart')

RTT plots: medians for 5-minute samples

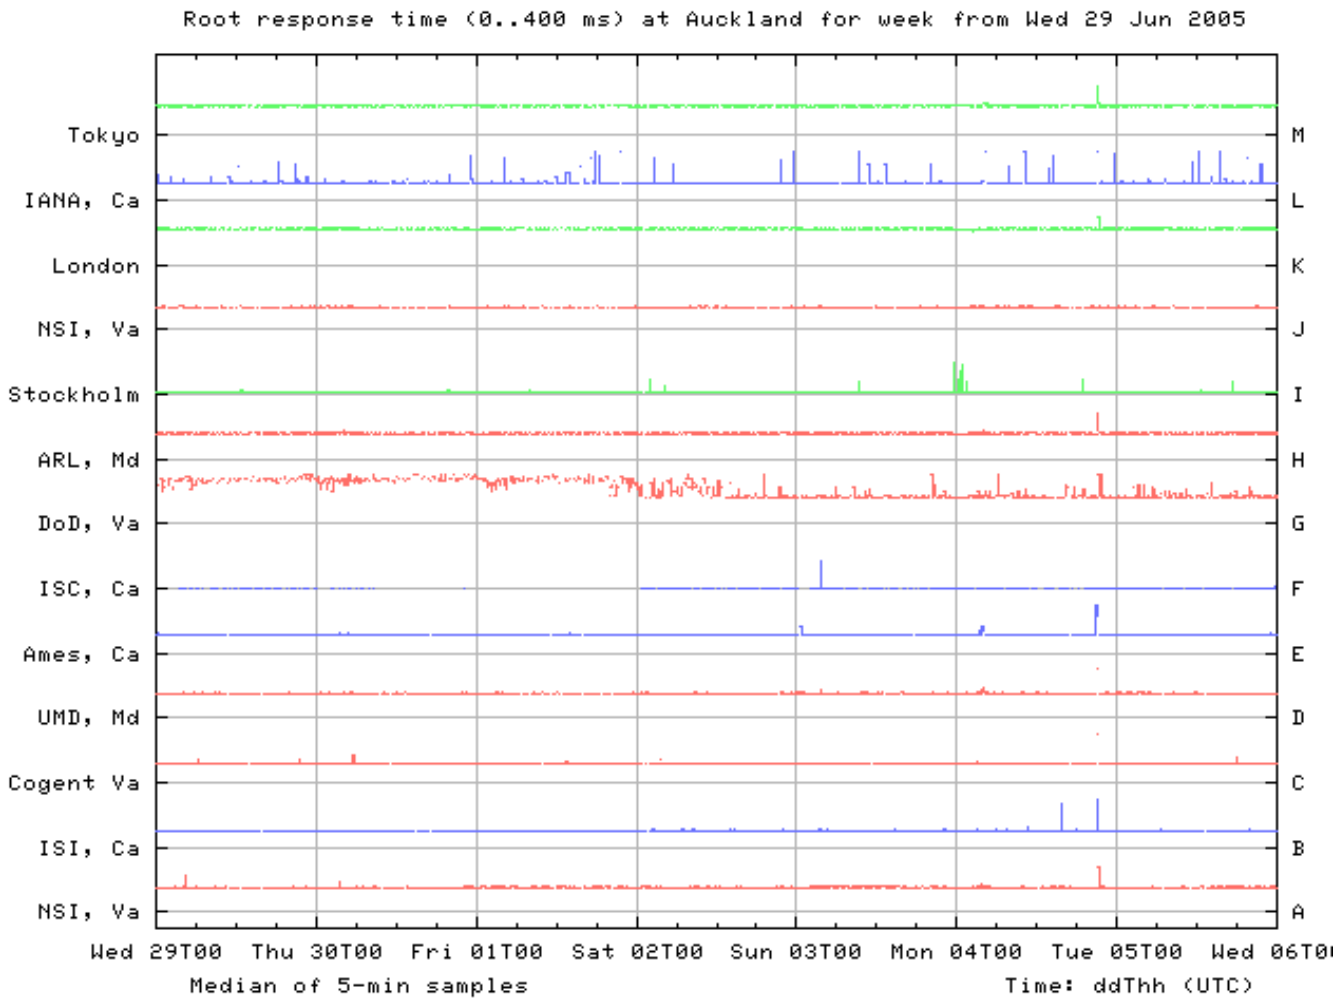

- **F and I roots are New Zaland anycast instances (see below)**
- **Median RTTs are mostly quite stable, except for L root**
- **G root has been overloaded for years —**
  note the improvement from Sat 2 July!

---

1. Example root RTT sequence plot: G root behaviour change

sequence plots for 5-minute samples

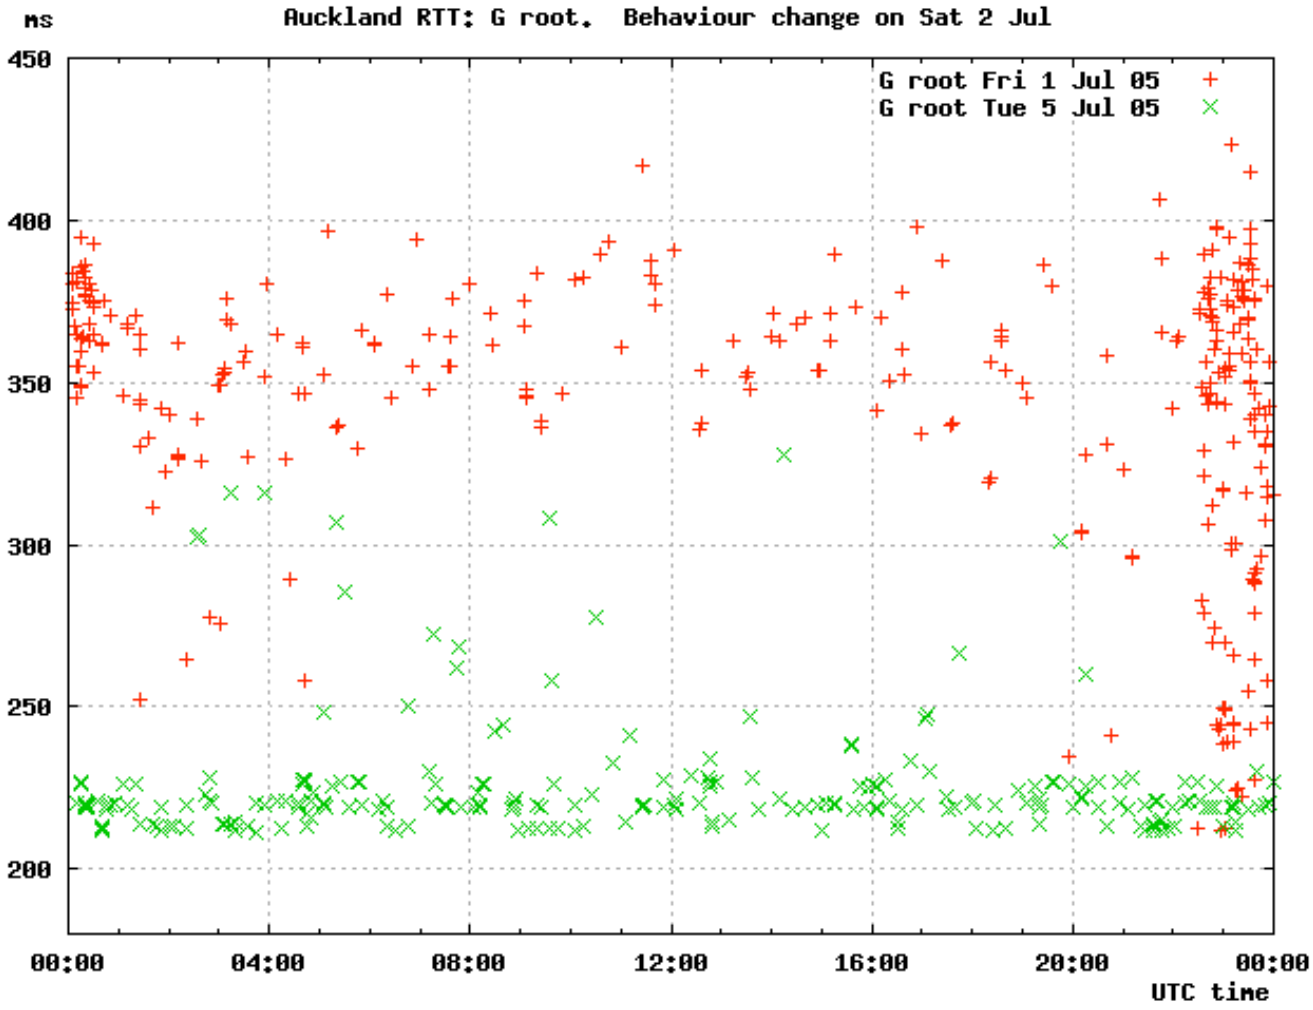

- **G root changed on 2 Jul 05**
  - ***RTT dropped from ~350 to 220 ms,***
  - ***Dispersion dropped***
- **Big improvement in G root performance ... thanks, G root folks :-)**
- **Observation: G now appears tri-modal from Auckland; we know it's not anycast!**

---

1. More sequence plots: anycast root RTT changes

- ***Can root RTT changes indicate responses from different anycast instances?***
- ***Root RTT changes are sometimes caused by server loading —***
  we see gradual dritfs in RTT, e.g. G root before 2 Jul 05
- ***RTT changes generally indicate route changes in paths to roots***
- ***Such route changes often last for several 5-minute intervals,***
  giving clear 'steps' in our strip charts

- ***Now that we have many anycast root instances, route changes could cause***

switvhing between different instances —
Can we see any evidence of that at Auckland or Boulder?

---

1. Auckland RTTs: New Zealand anycast roots

sequence plots for 5-minute samples

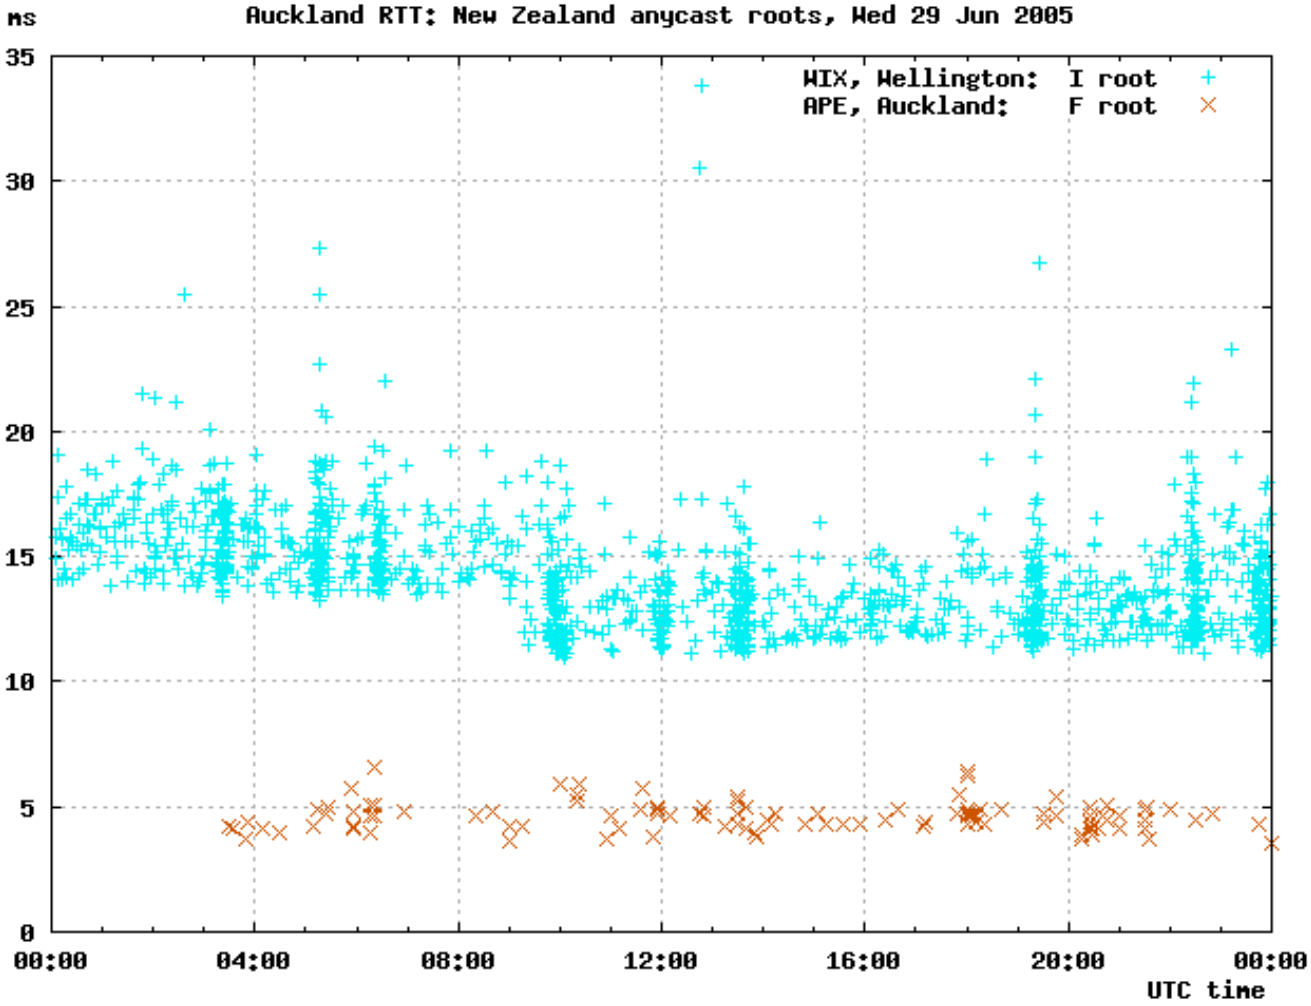

- **RTTs recorded in sequence without timestamps, plotted (above) at equal time spacings**
- **Higher RTT and more dispersion for I (Wellington) than F (Auckland)**
- **Drop in I's minimum RTT from 13 to 11 ms at 0930 must reflect a route change — there's only one instance of I in New Zealand!**

- ***[Puzzle: why does BIND prefer I to F, which is closer? Probably a routing policy artifact]***

---

1. Auckland RTTs for all anycast roots

sequence plots for 5-minute samples

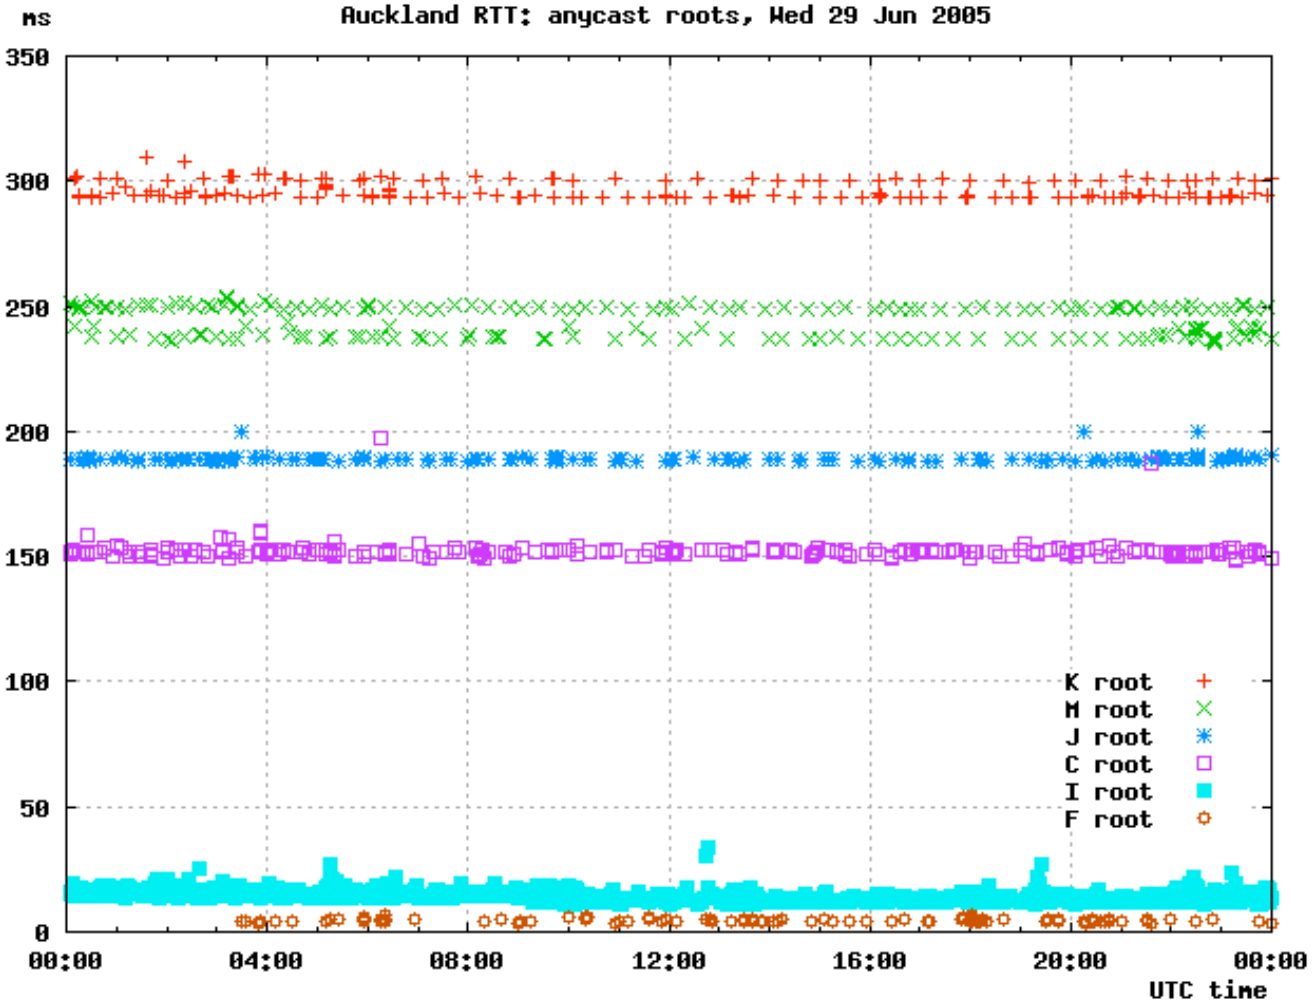

- **Roots outside New Zealand:**
  C in Los Angeles, J on US East Coast, M in Japan, K in Europe
- **K and M have two modes; are they different anycast instances?**
- ***Do we see bimodal behaviour at Colorado?***

---

1. Colorado RTTs for anycast roots

sequence plots for 5-minute samples

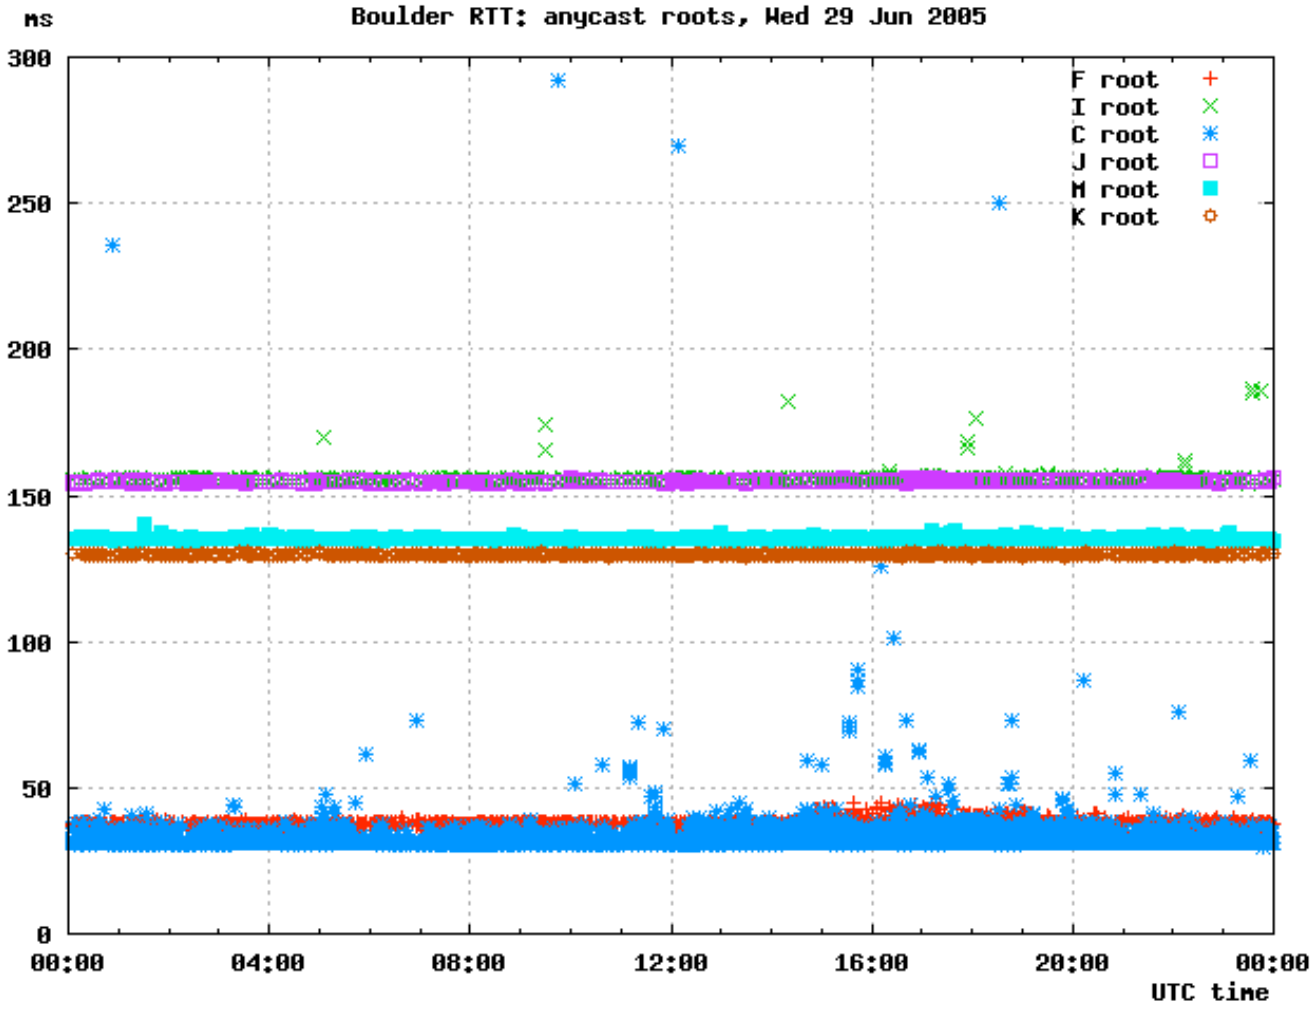

- **No sign of bimodal RTTs at Boulder**
- ***What did the Auckland RTTs look like before anycasting?***

---

1. Auckland RTTs for 'anycast' roots in 2002

sequence plots for 5-minute samples

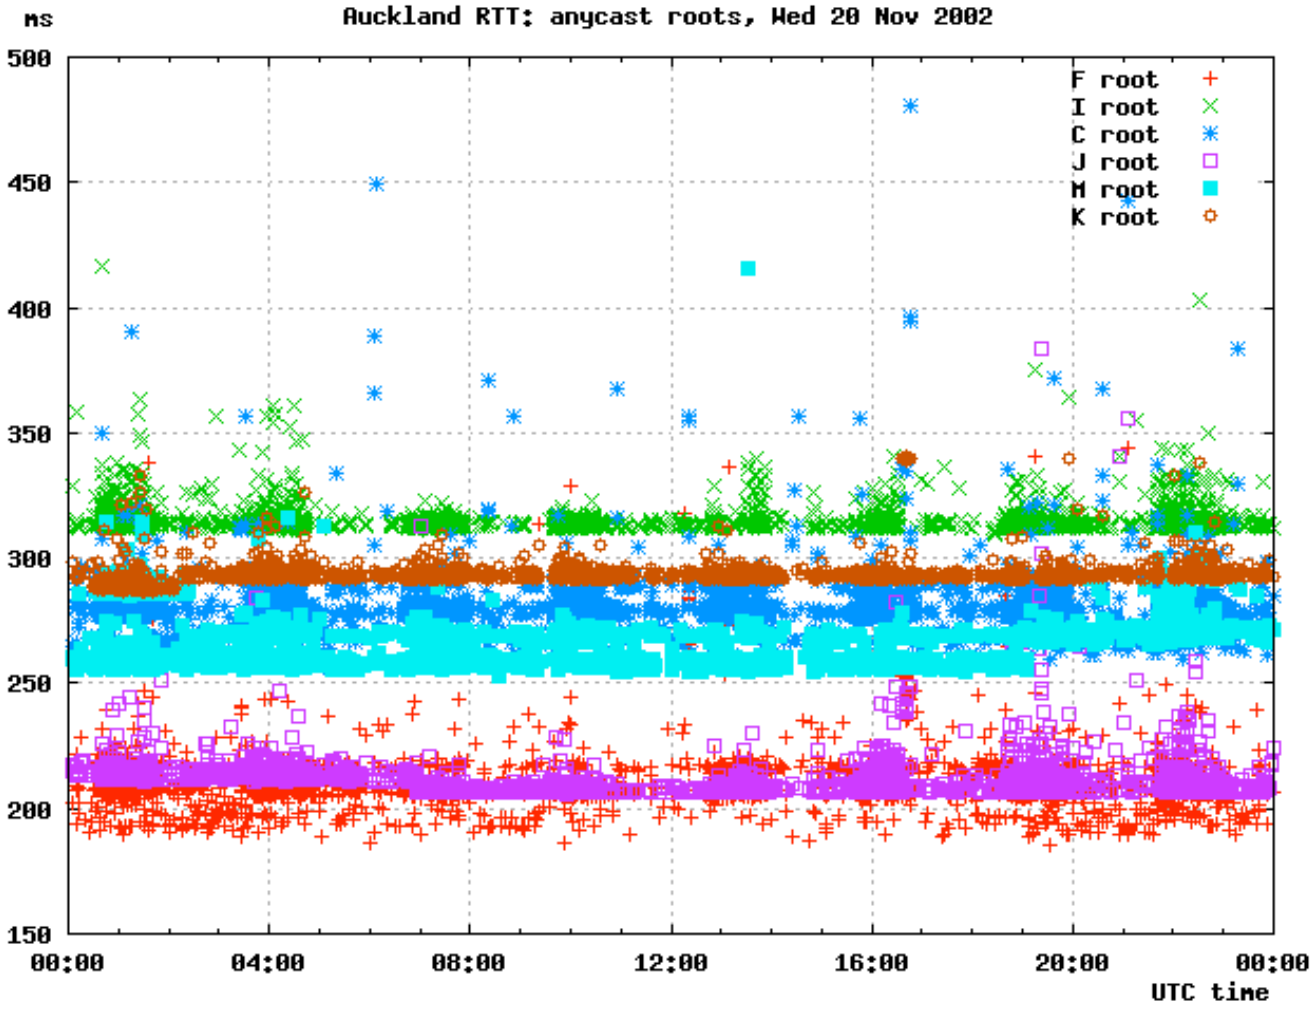

- **F, C and M appear bimodal**
- ***Let's have a closer look ...***

---

1. Auckland RTTs in 2002: F, C and M roots only

sequence plots for 5-minute samples

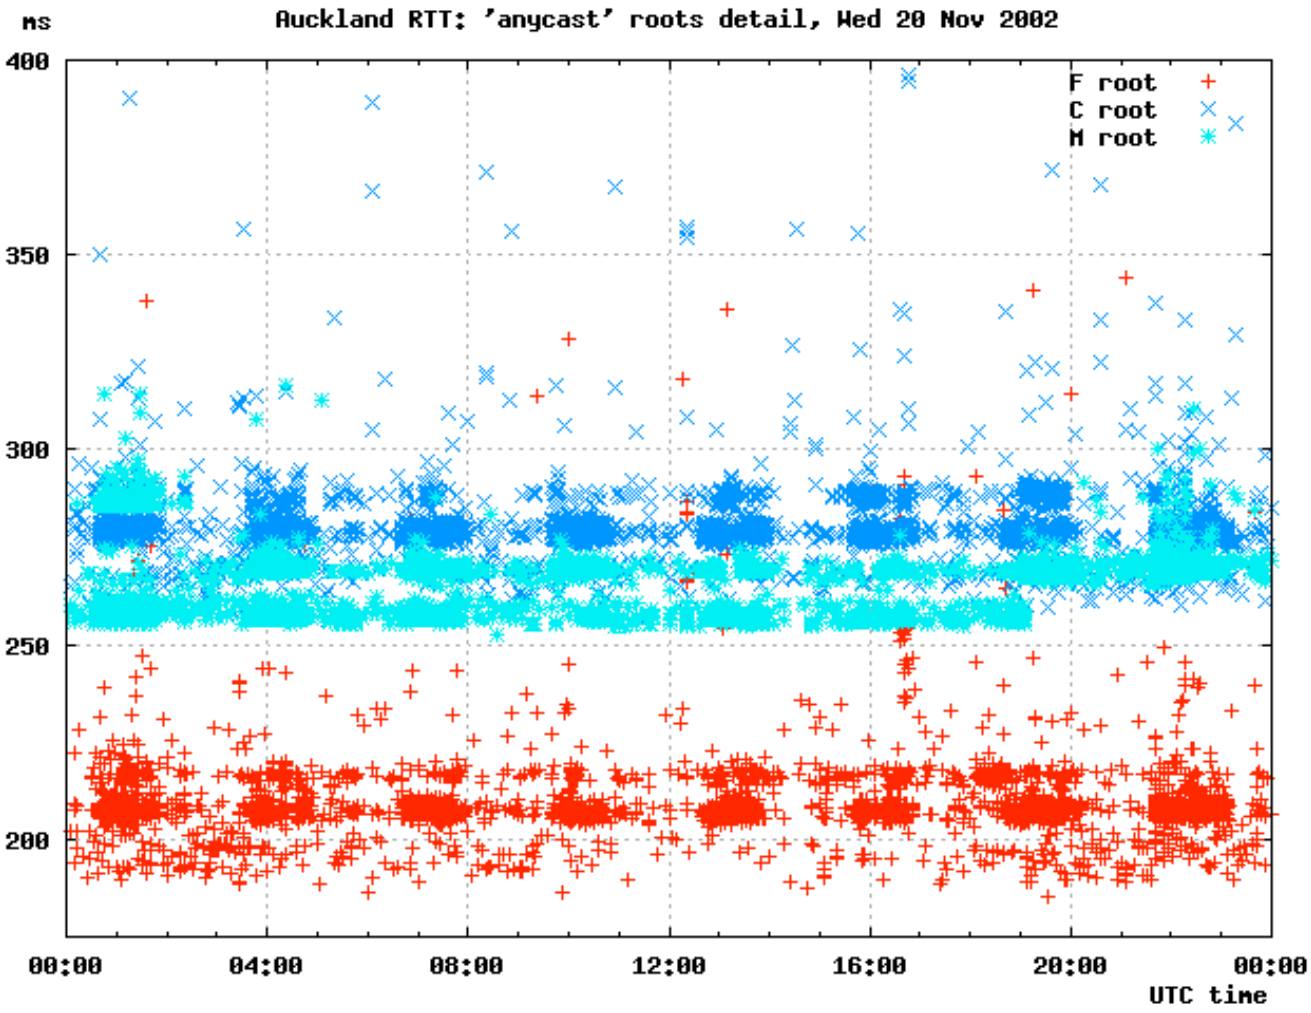

- **Lots of bimodal behaviour, as in Brownlee & Ziedins, PAM 2001**
- **But anycasting of roots didn't start until 2003 (??)**

---

1. Auckland RTTs in 2002: detail for F, C and M

sequence plots for 5-minute samples

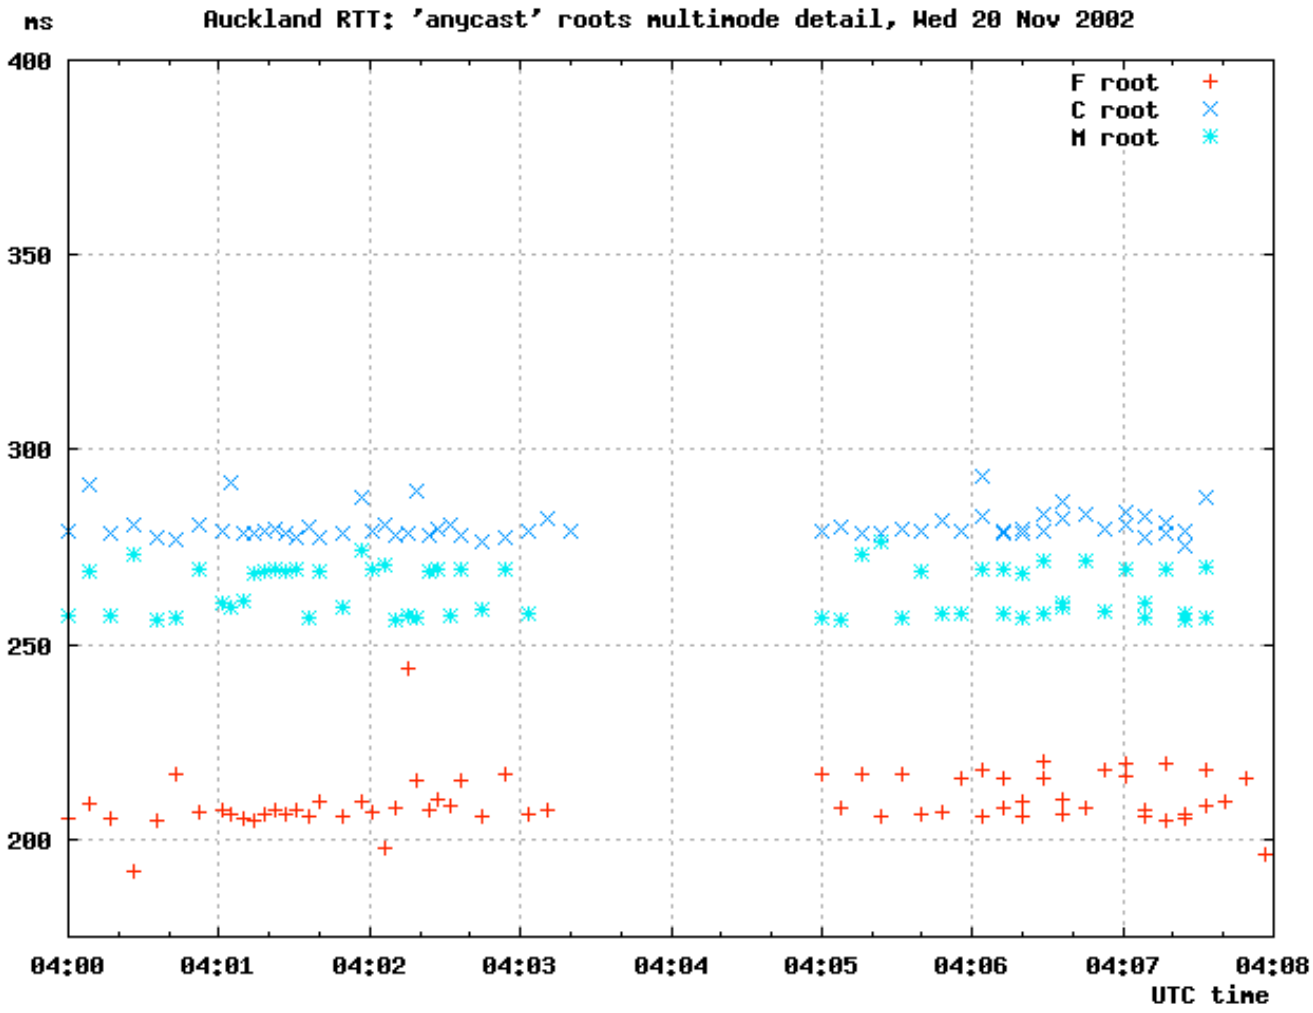

- **Sequence plots similar to those in Brownlee & Ziedins, PAM 2001**
- **That paper suggested that mutipathing was due to load balancing by some ISPs along the DNS packets' path; *that still seems a reasonable assertion***

---

1. Conclusion

- ***No evidence of route changes causing switching between anycast root instances at Auckland or Boulder***
- ***Big improvement in G root performance from 2 Jul 05***
- ***Much less variance in RTT distributions now (29 Jun 05, slide 7 above) compared to that in 2002 (slide 9 above) —***
  Internet backbone paths have become more stable!

Nevil Brownlee (nevil@caida.org)
Last updated: 20 July 2005

---

1. Extra slide: Colorado RTTs, week starting Sat 9 Jul 05

RTT plots: medians for 5-minute samples

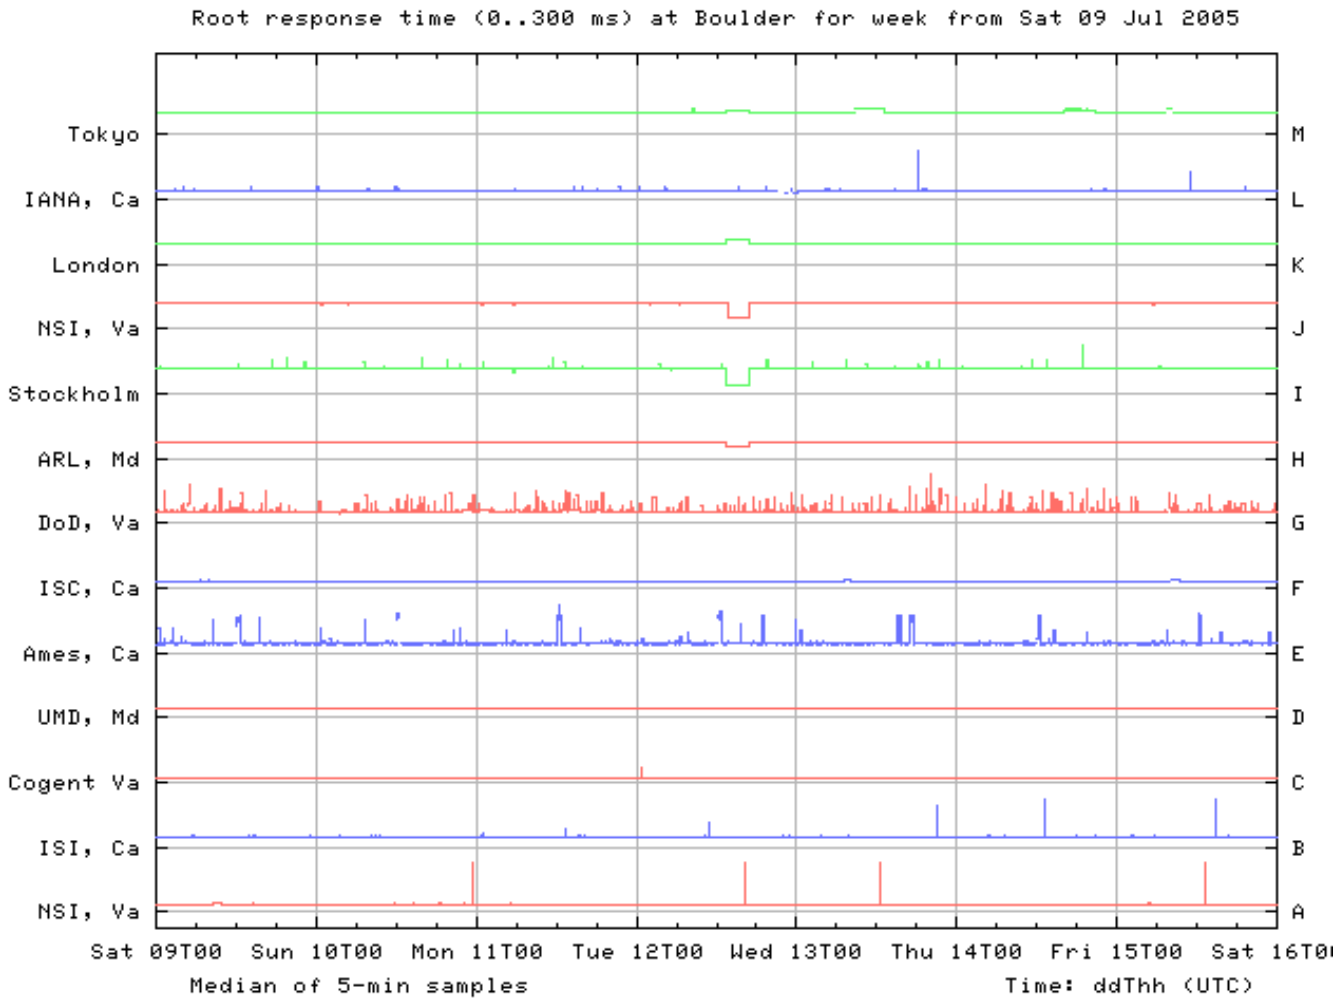

- **Current 'typical' week**
- **Most roots have stable RTTs**
- **H, I, J, K and M have long (about 6-hour) steps**
- **A, B, L have short spikes (one 5-minute interval)**
- **G no longer shows loading patterns, but has lots of spikes**
