DNS RTT Measurements: ccTLDs compared with roots/gTLDs

Nevil Brownlee
CAIDA, SDSC, UC San Diego and
The University of Auckland
nevil@caida.org

DNS, RIPE 45, Barcelona, May 2003

Overview

- Passive monitoring of DNS behaviour
- How DNS works
- NeTraMet setup to produce RTT charts
- Strip charts for roots/gTLDs, how to read them
- ccTLDs, NeTraMet setup to collect their RTT data
- RTT plots for ccTLDs
- Conclusion

Passive Measurements for Network Monitoring

- Active measurements (e.g. *ping, traceroute*) are easy to make. However, they perturb the network
- Passive measurements (e.g. *NeTraMet*) don't.
  Router-based measurements (e.g. *NetFlow, SNMP*) can affect router performance
- Passive measurements require a steady trickle of packets along the paths you're interested in
- DNS is a simple, easy-to-monitor service which is always present
- Nevil uses NeTraMet to observe DNS behaviour, and is looking for ways to use DNS for network monitoring

Global DNS behaviour

- User hosts send DNS requests to their local resolver
- If local resolver has requested domain in cache
  it responds directly
- Otherwise local resolver sends (non-recursive) request
  to a TLD, then walks down DNS tree
- When there are multiple NS records for a domain,
  local resolver distributes requests across all of them
  - BIND tries every server from time to time
  - BIND groups servers into sets and shares requests
    across all servers in the ‘closest’ set
  - Other nameserver implementations use different
    algorithms
- NeTraMet meter can see RTTs (request/response time)
  for all the root and gTLD servers
- Meter can also see RTTs for ccTLD servers, but there are
  fewer requests to them than to the gTLDs

NeTraMet setup at meter sites

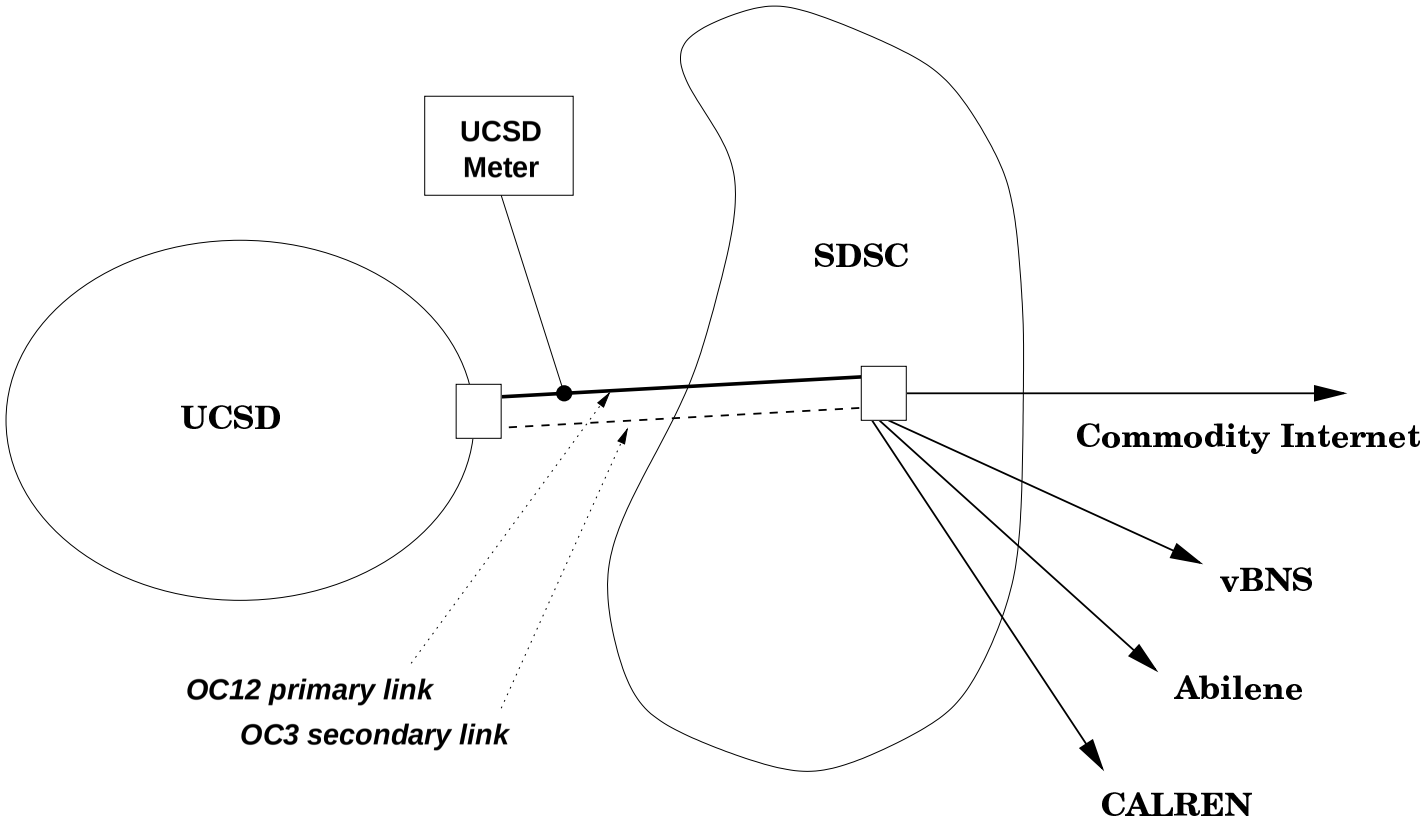

- CAIDA runs three meters now:
  - Auckland. 100Mb/s Ethernet via hub
  - Colorado. 100Mb/s Ethernet via SPAN port
  - UC San Diego. OC12 ATM via fibre splitter
- Current OC48 meter handles 150 kp/s
  (Dag 4 cards, 2 processors, multithreaded)

Root/gTLD performance web page

- www.caida.org/cgi-bin/dns_perf/main.pl
- Passive obervations from Auckland (*ua*), Colorado (*cu*) and San Diego (*ucsd*)
- Data from early January 2003 (all sites, UCSD from late 2001)
- RTT (ms) loss% and count for all sites
- Web page lets you select the days / sites / metrics of interest
- Plots show 5-minute medians, scaled to fit 3/4 of space between server lines

Example Strip Charts: root RTTs, UCSD, 22 March 03

Show ☑ rtt ☐ losspc ☐ count strip charts for ☑ root ☐ gTLD servers observed from ☑ ucsd ☐ ua ☐ cu ☐ sjc

Start date (UTC): 2003 / 3 / 22 for 7 days < | Plot >

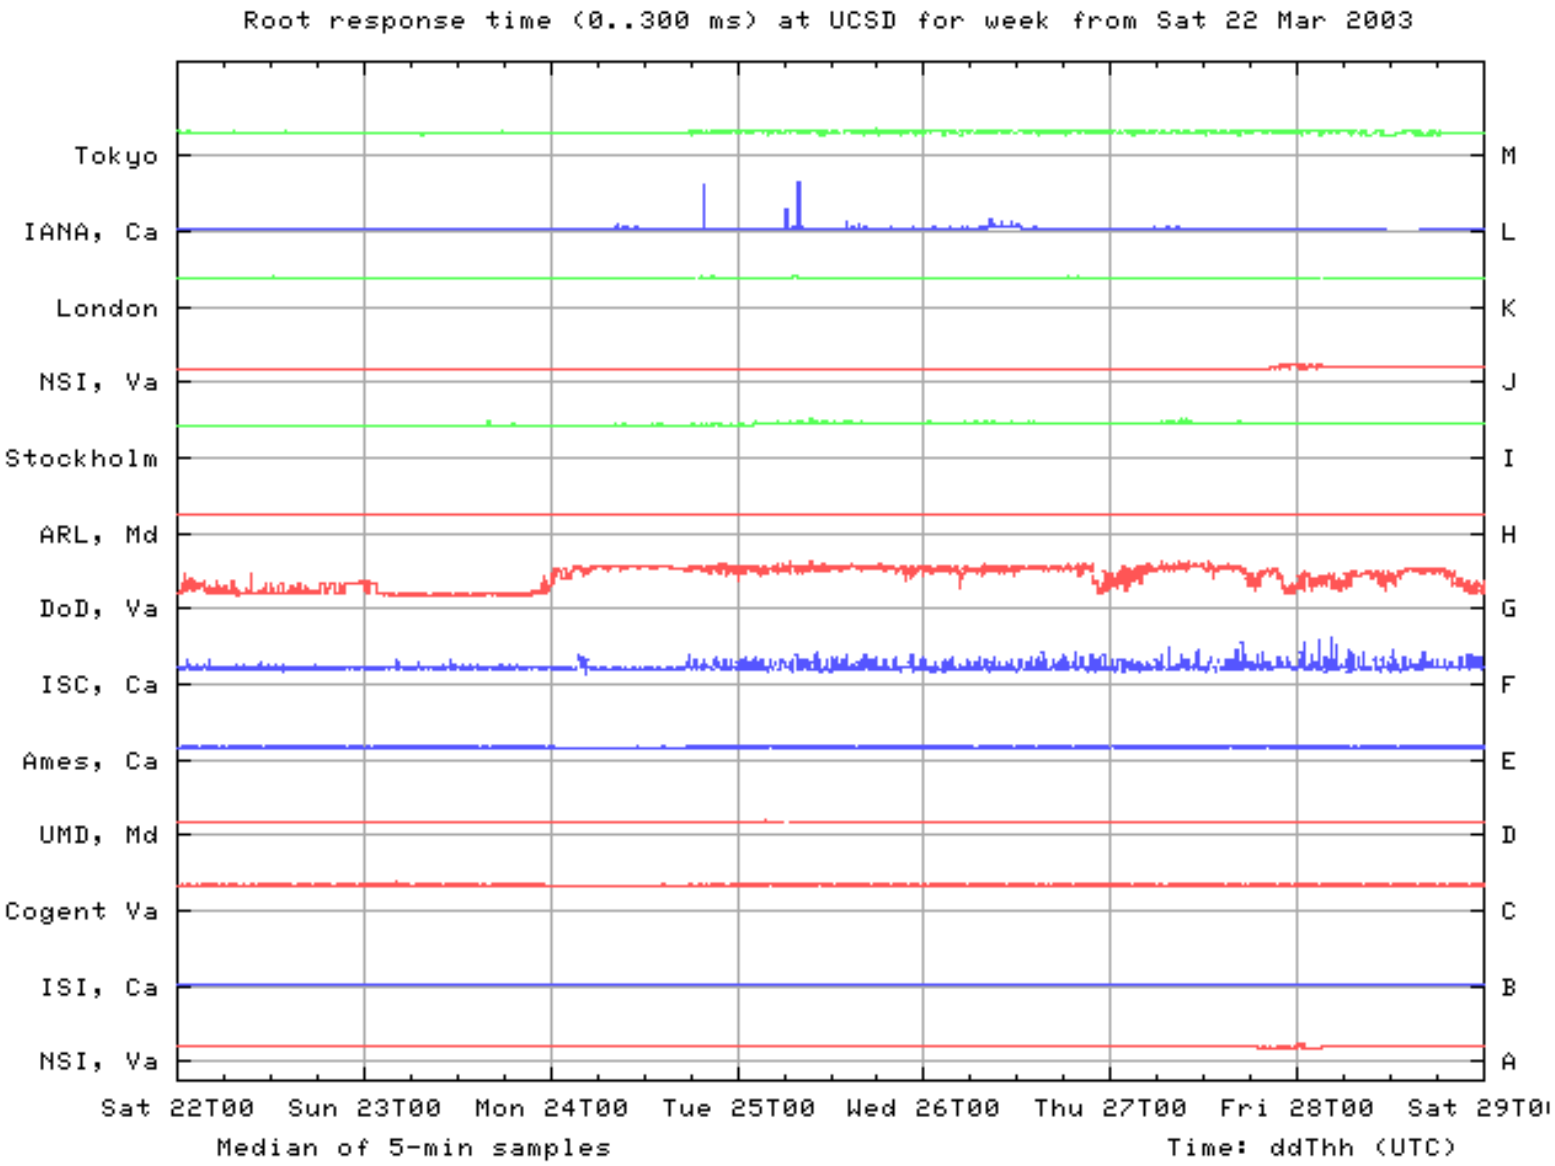

Roots from UCSD, 19 April 03

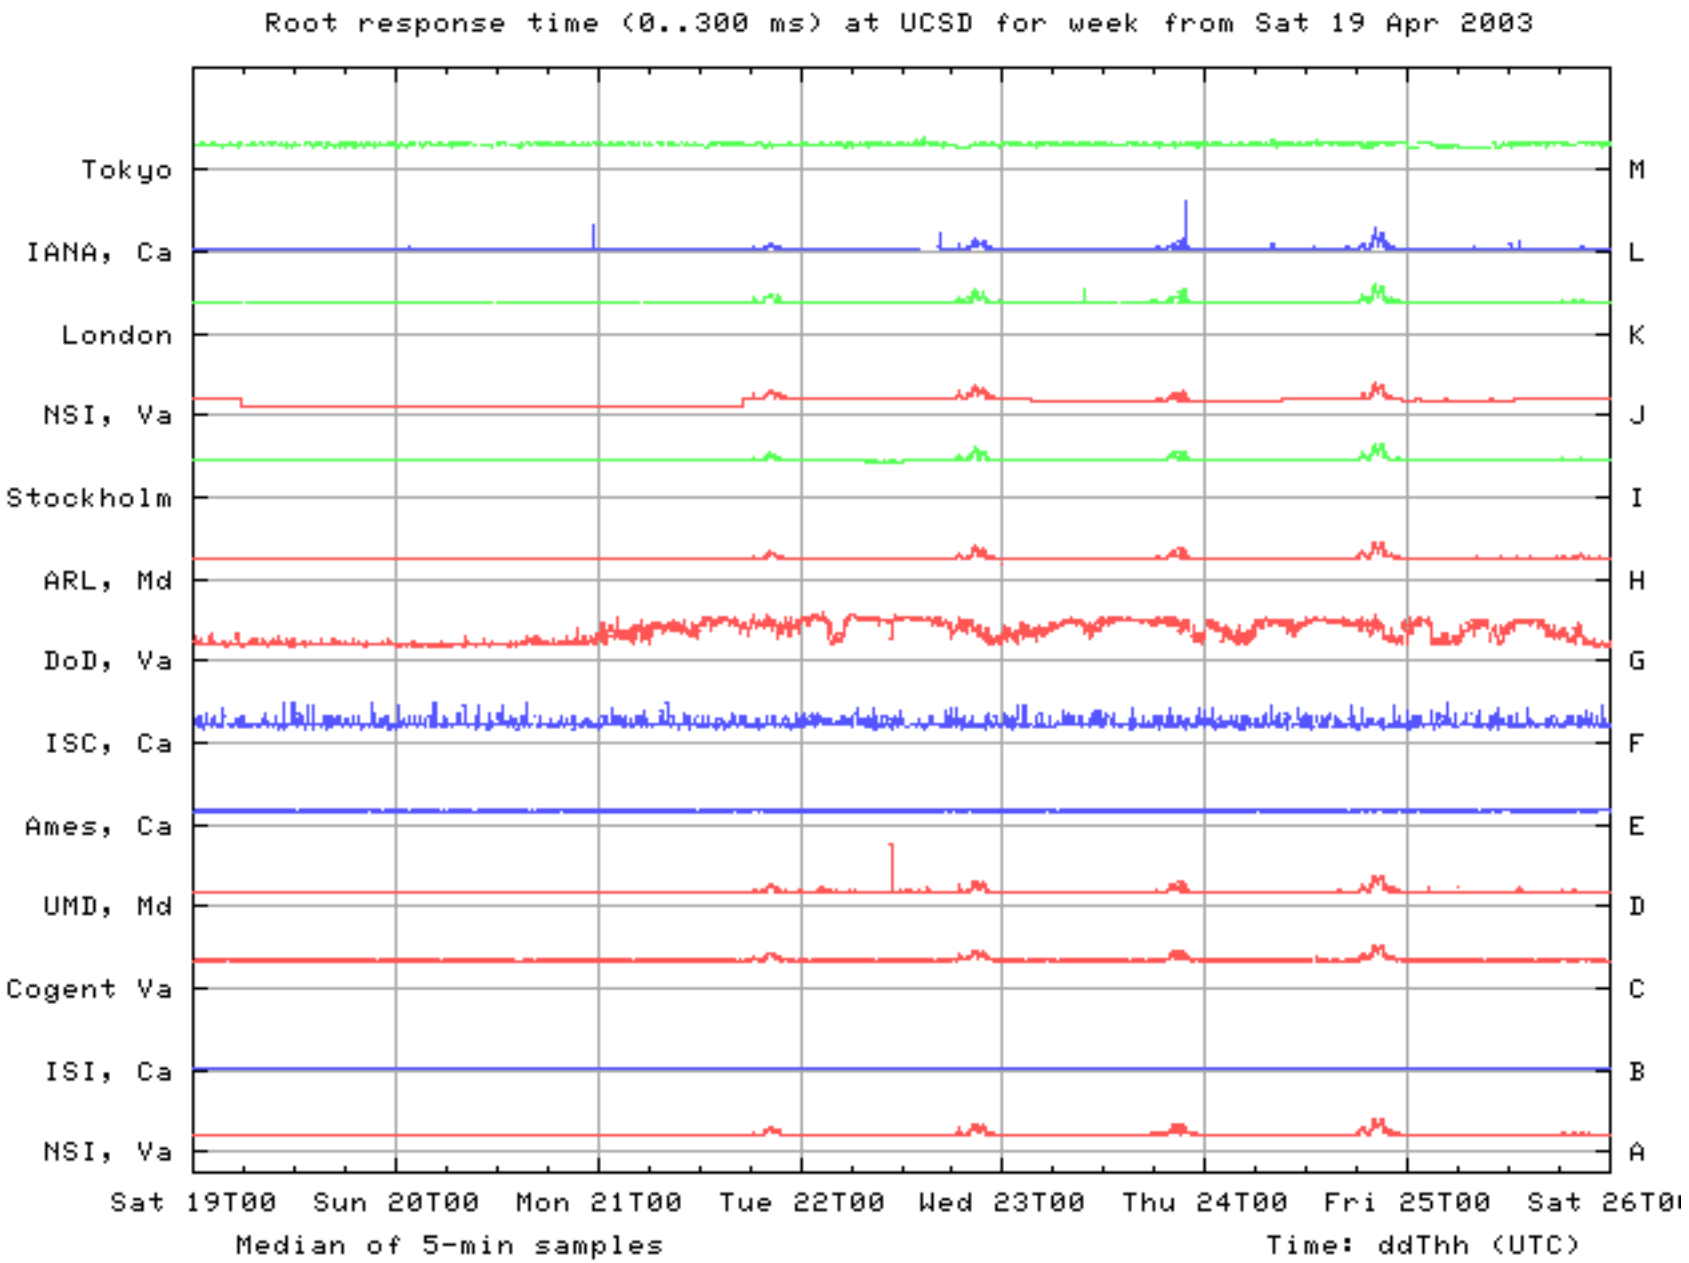

- Daily drifts in G show congestion of G root server
- Weekday bumps Monday-Thursday: Congestion at UCSD, 2000-2200 UTC (1100-1300 PDT).

gTLDs from Auckland, 19 April 03

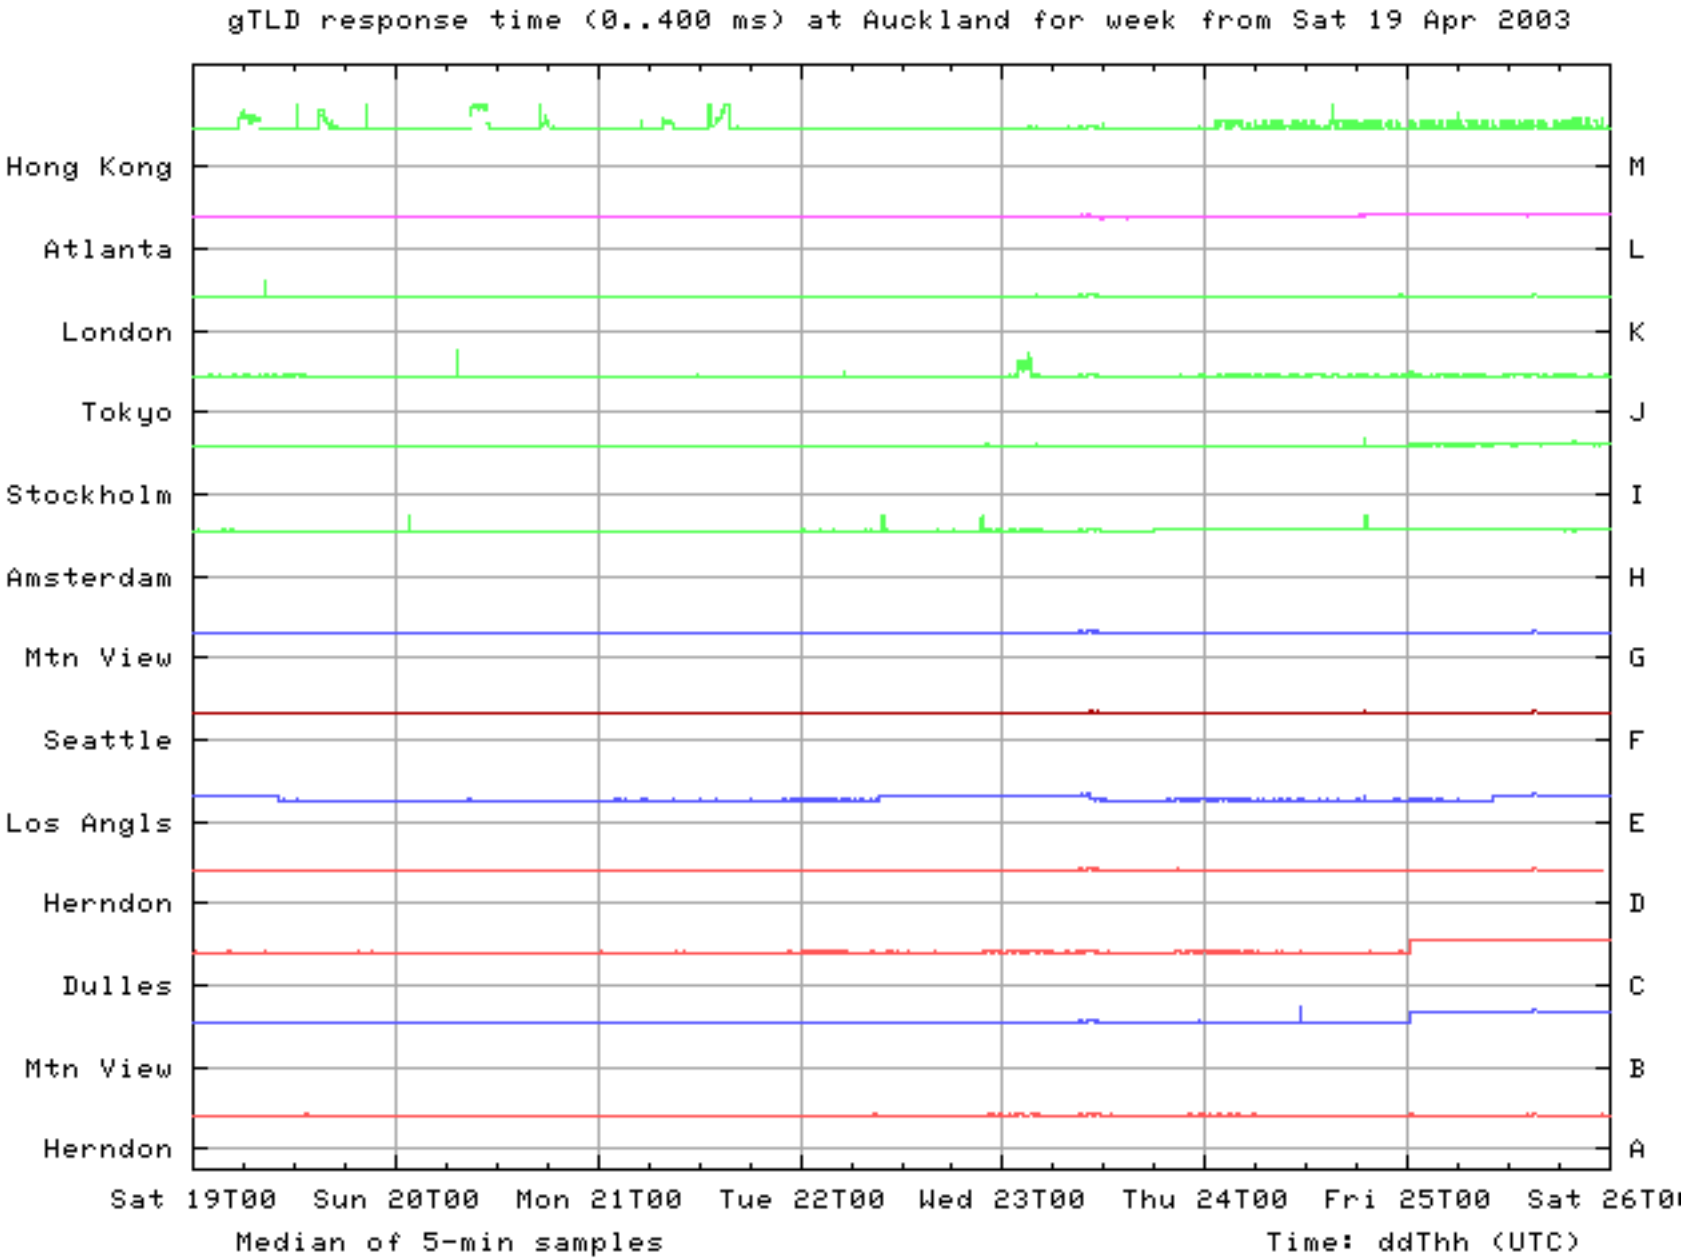

- Steps in E, B+C ⇒ route changes near server
- Blips in A-M Wednesday 1000 ⇒ route changes close to meter

What do RTT strip charts show?

- Effects caused by route changes (*steps*),
  server loading and network congestion (*drifts* and *bumps*)
  - Changes near server $\Rightarrow$ affects only that server
  - Changes near site $\Rightarrow$ affects lots of servers
- Monitor
  - Local resolver behaviour (verify that resolver is working!)
  - Connectivity to sites of interest (Auckland - US) e.g. to gTLDs or specified ccTLDs
  - Congestion on links to Internet
  - Global DNS behaviour

ccTLDs: How do they differ from root/gTLDs?

- Only 13 roots and 13 gTLDs, small set of fixed addresses
  [Need to see effect(s) of root server anycasting]
- To build NeTraMet ruleset for ccTLDs:
  - Get list of country codes from web
  - Make dig input file to query the country code domains
  - Run dig to get cc domain info from a root server
- There are 237 country codes, with varying numbers of nameservers
- Servers can be:
  - *Own:* server in its own ccTLD,
  - *Other:* server in some other TLD, or
  - *Multi:* server for many ccTLDs
    - Full ruleset has 106 multi-servers
    - NS.RIPE.NET serves 75 country code domains

Country Code distribution at UCSD

191 countries by nbr of rtts (rho):

	rtts	%	ccd%	trans	%	ccd%	servers			Country
0	929008	(37.9,	62.1)	1236038	(33.0,	67.0)	108	0	??	Multi (88 count
1	324189	(13.2,	48.9)	393417	(10.5,	56.4)	8	201	sr	Suriname
2	136172	(5.5,	43.4)	173025	(4.6,	51.8)	6	113	kr	Korea, Republic
3	100690	(4.1,	39.3)	137405	(3.7,	48.1)	4	73	fi	Finland
4	92542	(3.8,	35.5)	109144	(2.9,	45.2)	6	108	jp	Japan
5	86506	(3.5,	32.0)	118748	(3.2,	42.1)	4	177	ru	Russian Federat
6	66791	(2.7,	29.3)	119029	(3.2,	38.9)	10	81	de	Germany
7	53637	(2.2,	27.1)	66488	(1.8,	37.1)	3	224	us	United States
8	47541	(1.9,	25.1)	69094	(1.8,	35.2)	5	38	ca	Canada
9	44032	(1.8,	23.4)	60018	(1.6,	33.6)	5	106	it	Italy
10	38554	(1.6,	21.8)	47992	(1.3,	32.4)	2	190	sl	Sierra Leone
11	37575	(1.5,	20.2)	40532	(1.1,	31.3)	1	169	ph	Philippines
12	25158	(1.0,	19.2)	33344	(0.9,	30.4)	3	162	pk	Pakistan
13	24667	(1.0,	18.2)	109308	(2.9,	27.5)	2	225	um	US minor islanc
14	23445	(1.0,	17.3)	29145	(0.8,	26.7)	3	160	no	Norway
15	21445	(0.9,	16.4)	29523	(0.8,	25.9)	5	47	co	Colombia
16	20965	(0.9,	15.5)	22035	(0.6,	25.3)	2	30	br	Brazil
17	18277	(0.7,	14.8)	42724	(1.1,	24.2)	4	145	mz	Mozambique
18	16413	(0.7,	14.1)	38076	(1.0,	23.1)	5	194	sb	Solomon Islands
19	16126	(0.7,	13.5)	40849	(1.1,	22.1)	9	171	pl	Poland
20	15579	(0.6,	12.8)	27423	(0.7,	21.3)	8	21	be	Belgium

Country Codes at three sites

	San Diego	Colorado	Auckland	
CCs	191	179	182	
multi	0, 37.9%	0, 47.5%	0, 26.0%	(ix, %observations)
sr	1, 13.2%	1, 10.3%	2 , 5.7%	
kr	2, 5.5%	12, 1.2%	4, 2.6%	
ru	5, 3.5%	3, 3.7%	5, 2.3%	
de	6, 2.7%	4, 2.2%	6, 2.3%	
jp	4, 3.8%	5, 1.9%	7, 1.8%	
ca	8, 1.9%	7, 1.6%	8, 1.7%	
it	9, 1.8%	9, 1.3%	13, 0.9%	
no	14, 1.0%	28, 0.4%	28, 0.4%	
br	16, 0.9%	17, 0.8%	26, 0.5%	
nz	27, 0.5%	25, 0.5%	1, 28.8%	

- Similar numbers (182 to 191) of ccTLDs observed at three meter sites
- 26 to 47 % of ccTLD lookups are from multi-servers: they are (nearly) as important as TLDs
- Multi-servers 1,6,10,36 are in common used at all three sites

Multi-Servers from three sites

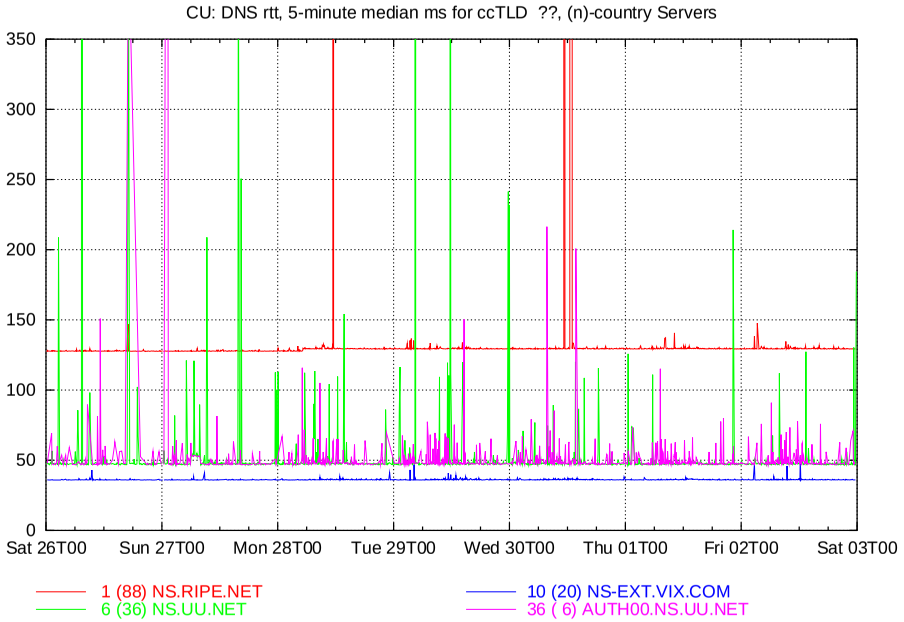

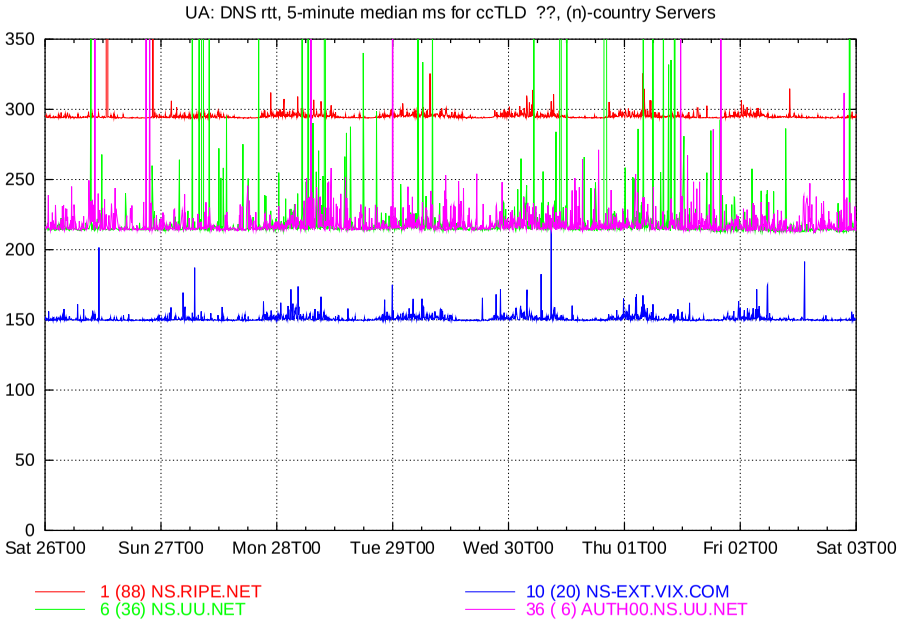

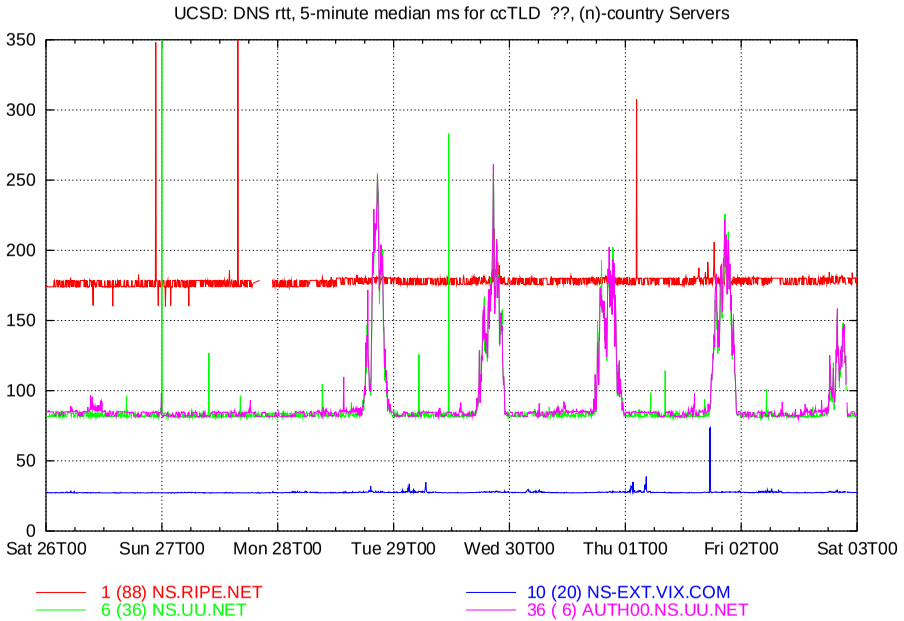

- Only Multi servers have enough data to plot for *all three* sites
- UA daily congestion0600-1000 (red and blue), UU.NET 2200-1000 (green and magenta)
- UCSD daily congestion 2000-2200 (green and magenta)
- RIPE.NET (red) and VIX.COM (blue) are stable
- UU.NET is comparatively noisy
- Overall Multi servers are (nearly) as stable as root servers

Busy ccTLDs (1): sr, Suriname

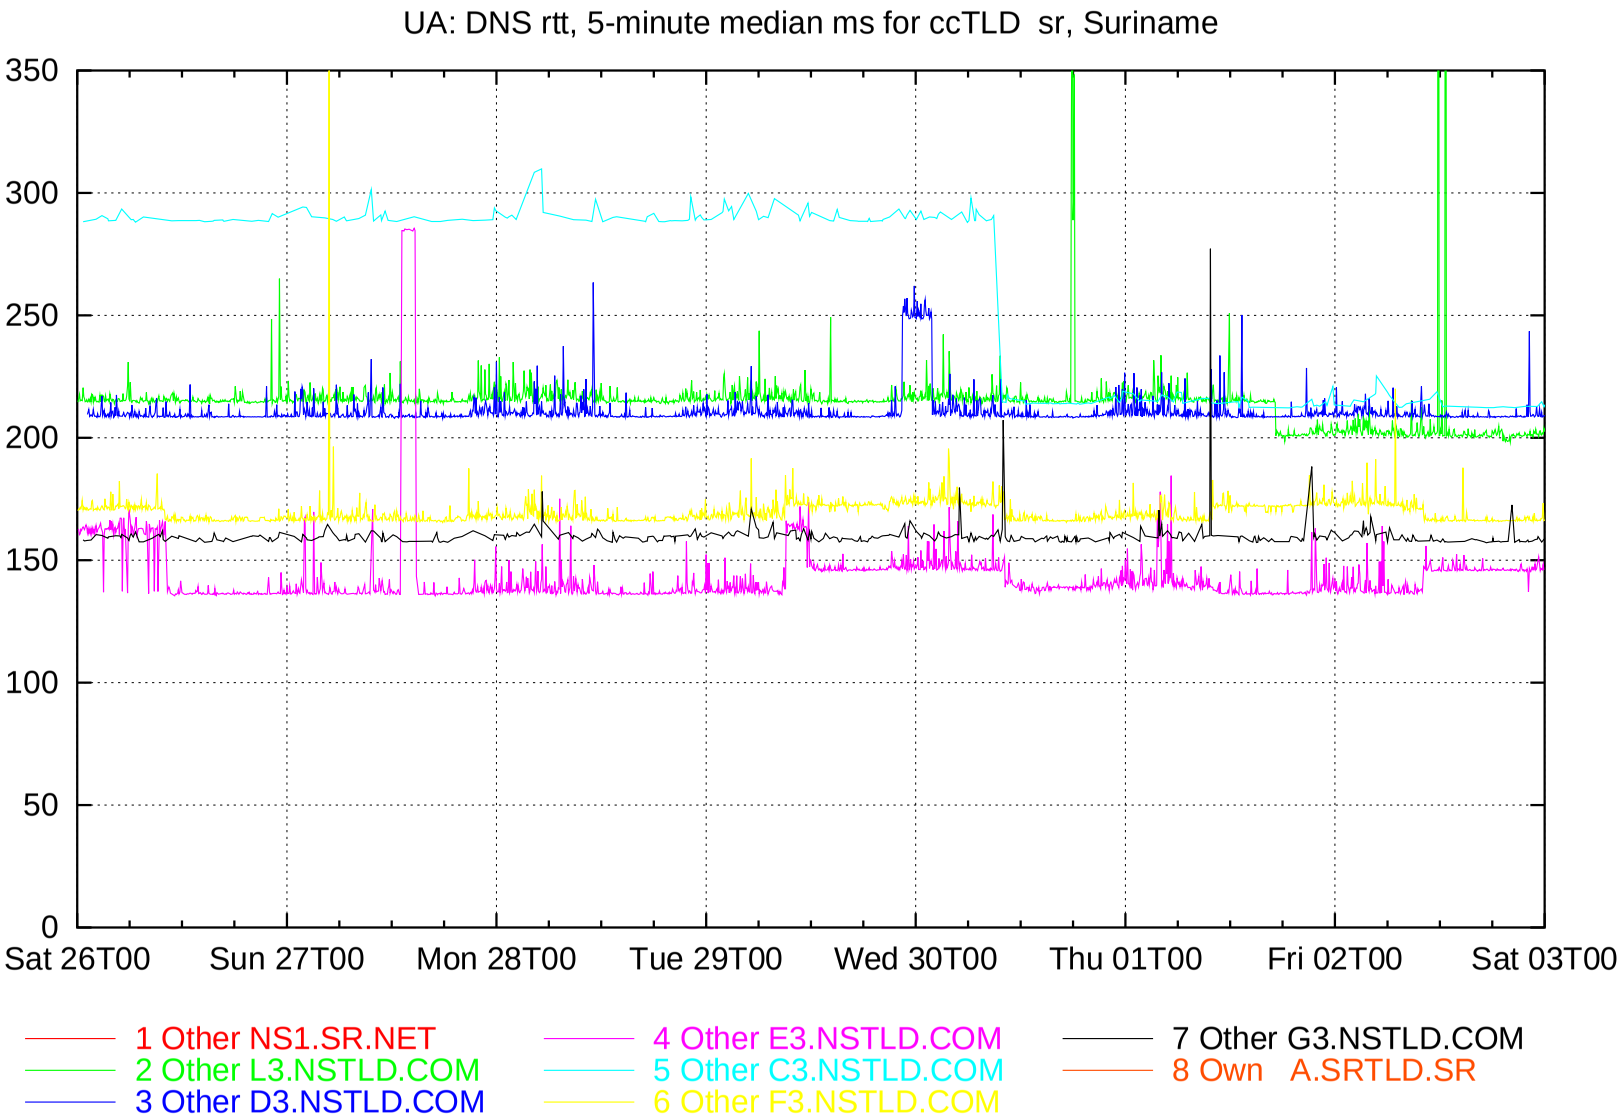

- Seven servers observed, behaviours differ
- Daily load congestion 2200-1000, *not the Auckland weekday load time of 0600-1000*
- Does anyone know what's in Suriname to attract DNS traffic?

Busy ccTLDs (2): de, Germany

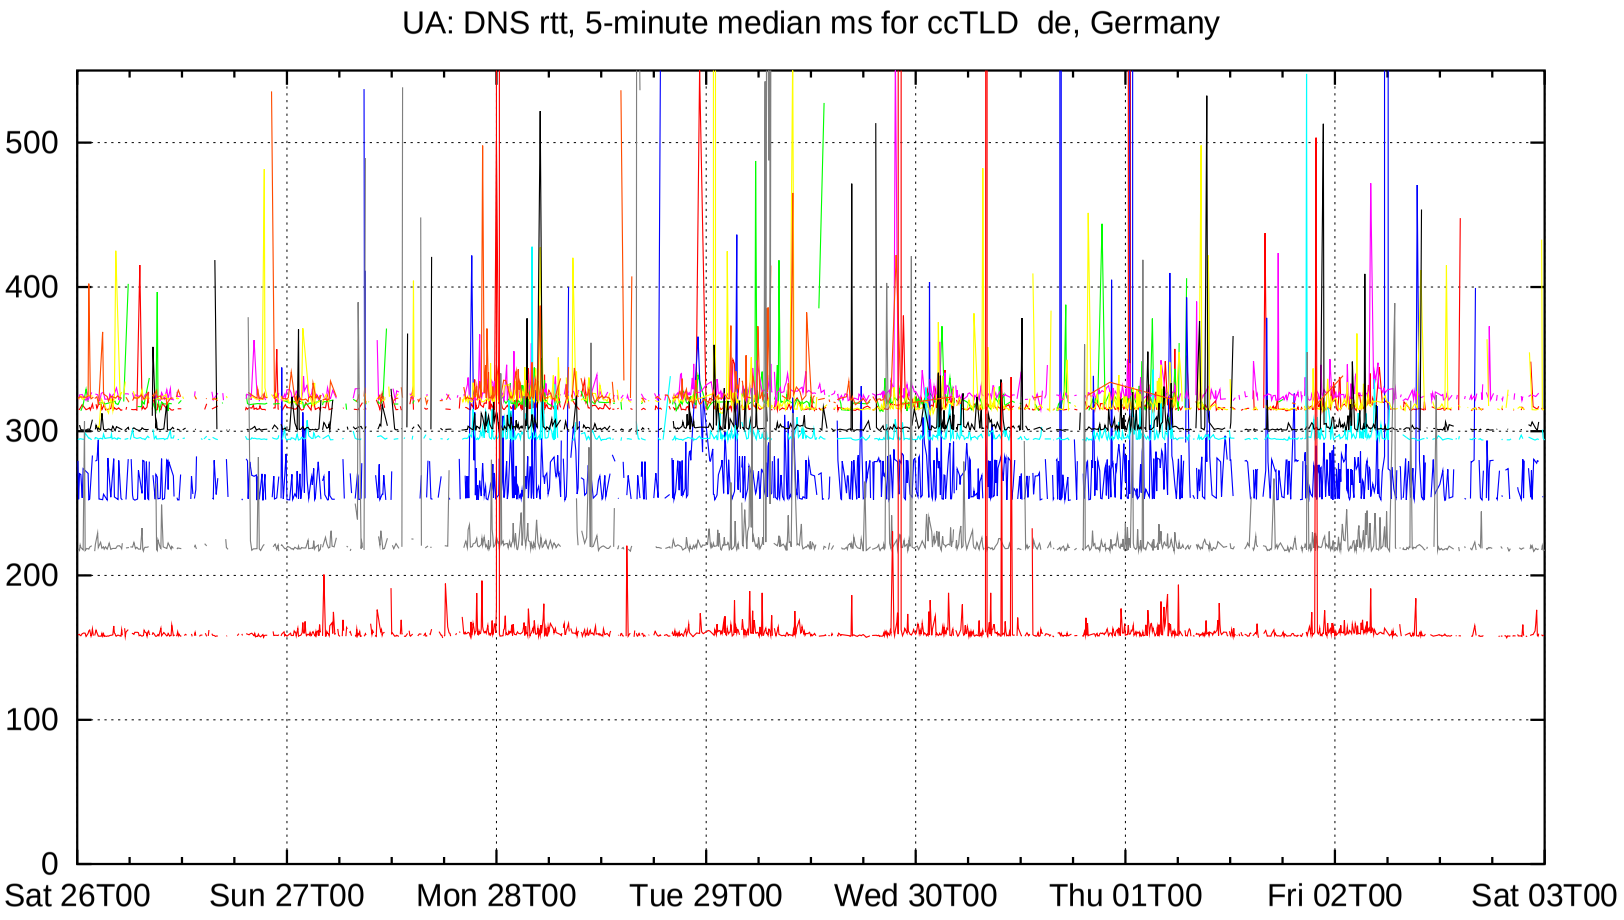

- Seven servers observed, behaviours are similar
- Daily load congestion 2200-1000, same as for .sr
- Variation for .jp (server 8) higher than for other servers

Busy ccTLDs (3): it, Italy

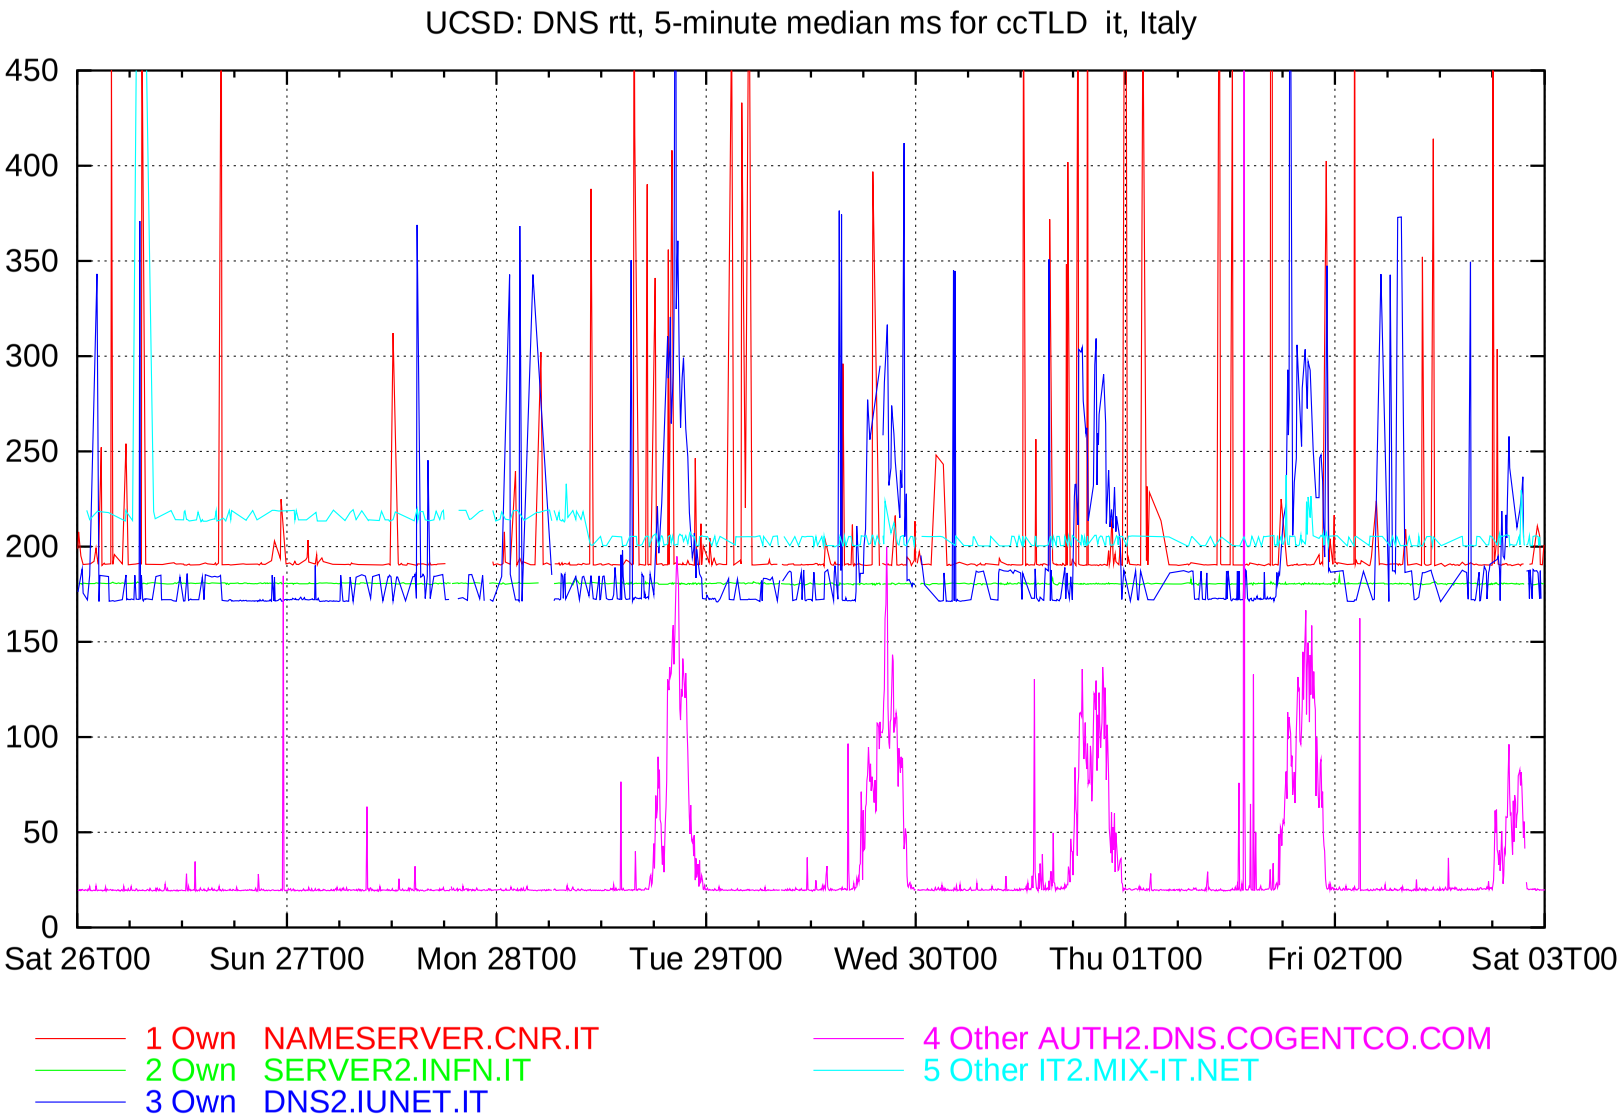

- Five servers observed, behaviours differ
- UCSD weekday load congestion 1800-2300 - server in US gives higher rtt!
- Servers in .it have very different variation patterns

Busy ccTLDs (4): br, Brazil

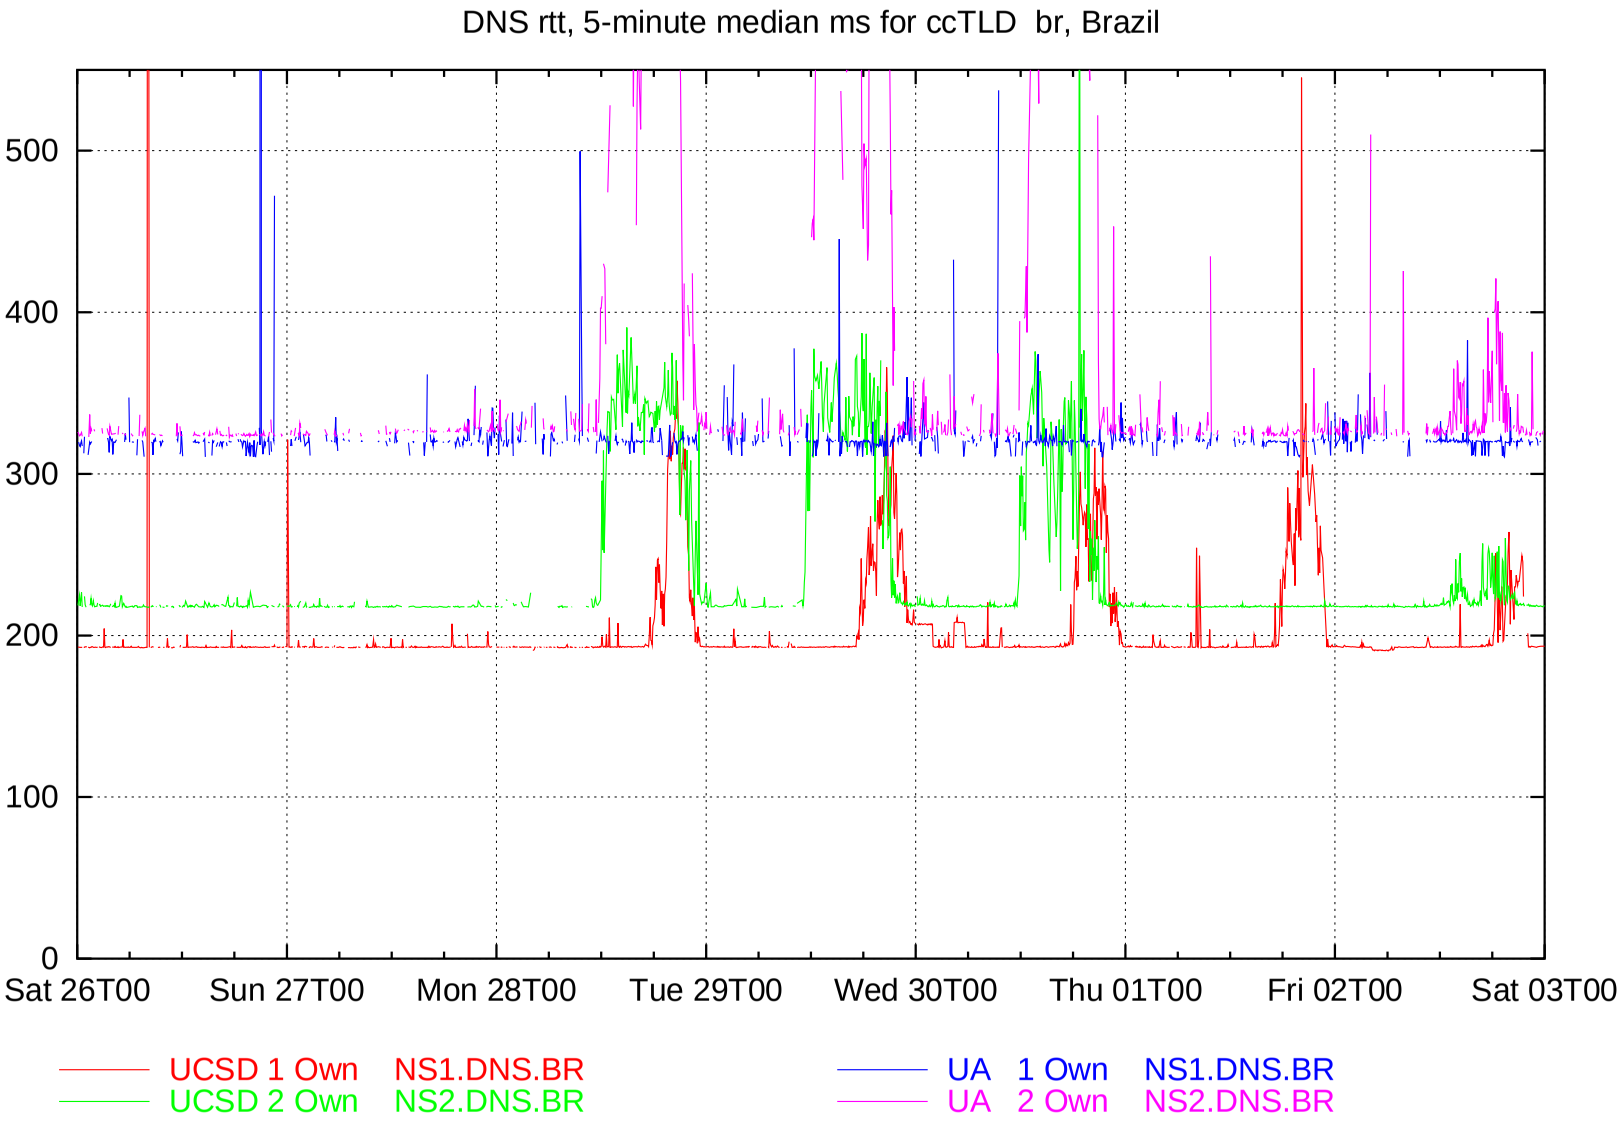

- NS1 (red trace) shows UCSD weekday loading
- NS2 (green trace) shows congestion 1100-2200 ?
- UA traces are similar (blue and magenta), but UA NS1 trace doesn't have a dominant minimum ??

Busy ccTLDs (5): Daily Congestion plots

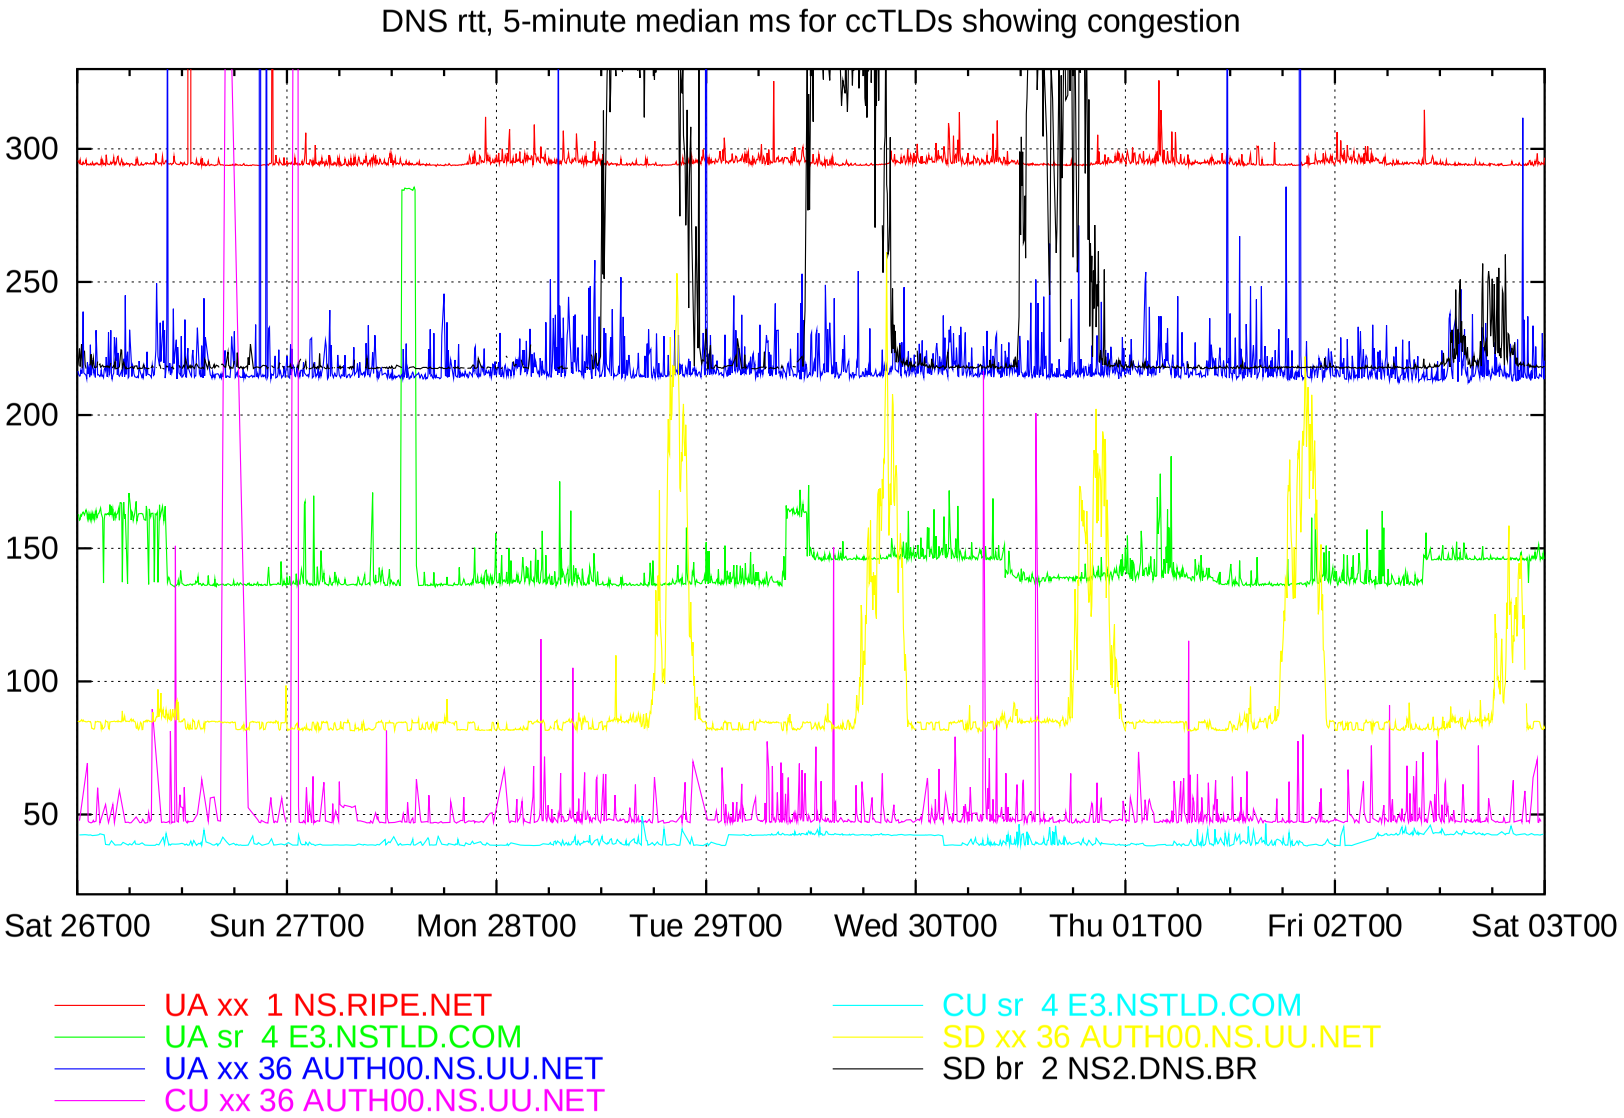

- UA (red, green) are similer
- CU (magenta, cyan) are different, cyan is later than magenta
- SD (yellow, black) are different. Yellow as for roots, black lasts for much longer periods
- Congestion seems to be *mostly* caused by local link,

Conclusion

- Strip charts are useful to verify local resolver is working, and to monitor connectivity to TLD and ccTLD nameservers
- Need to make a ‘user package’ which sites could use to produce their own strip chart web pages
- Also need to make software to automatically spot changes in the strip charts
- ccTLDs could be useful for monitoring
  - There are many more of them than roots/gTLDs, they’re more geographically distributed
  - But they don’t get as many lookups as the roots/gTLDs. Maybe the ones your users find most popular would (?)
- Nevil needs some more NeTraMet meter sites for this DNS monitoring work, see web page

  `http://www.caida.org/~nevil > Setting up a NeTraMet meter`