Bgpdump2: A Tool for Full BGP Route Comparison

Yasuhiro Ohara
yasuhiro.ohara@ntt.com
NTT Communications

20150319
p19 NTT

heatmap/oregon-ix2-rib.20150319.0000-p19

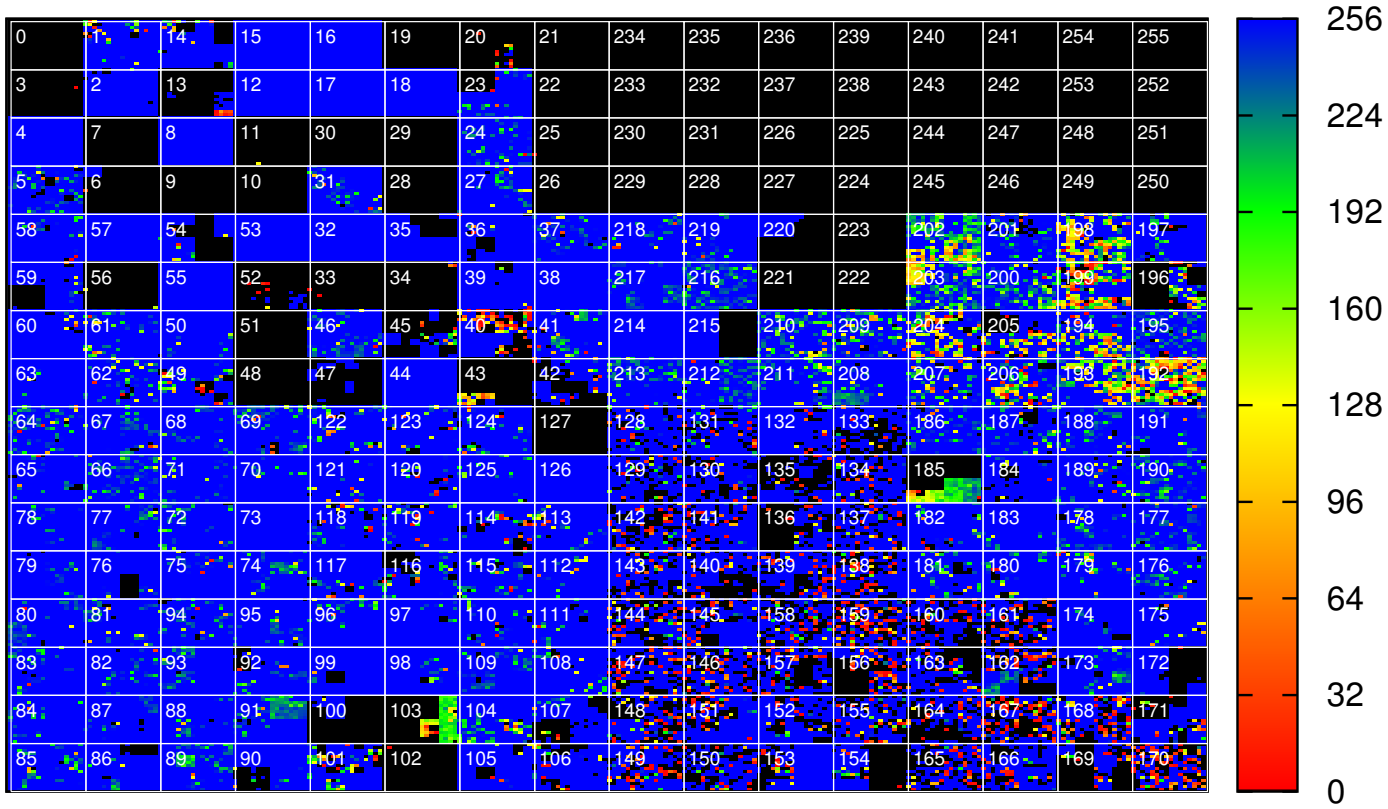

Motivation (1)

- As a researcher in an ISP ...
- Don't we need a way to evaluate a shape of (our) BGP full route routing table ?
- Why ?
  - Mistakes: missing routes, route leaks, ...
  - Better future configuration: roundabout routes
- Many things boils down to it
  - the routes are good, then we're good.

Motivation (2)

- How can we evaluate a shape of BGP full routes ?
  - We don't know it yet.
  - detailed (statistical) analysis ?
  - compared to others ?
- So let's create a tool to help doing those.

bgpdump2 summary

- Full scratch in C (4000- lines)
- Open-source software
  - <https://github.com/yasuhiro-ohara-ntt/bgpdump2>
- Capability:
  - supported: bz2, gzip, and raw, MRT TABLE_DUMP_V2 format, ipv4 / ipv6 routes
  - show statistics per peers (e.g., #routes, #nexthops, #unique-AS-paths)
  - routing table construction, longest-match table lookup

route collectors and monitors

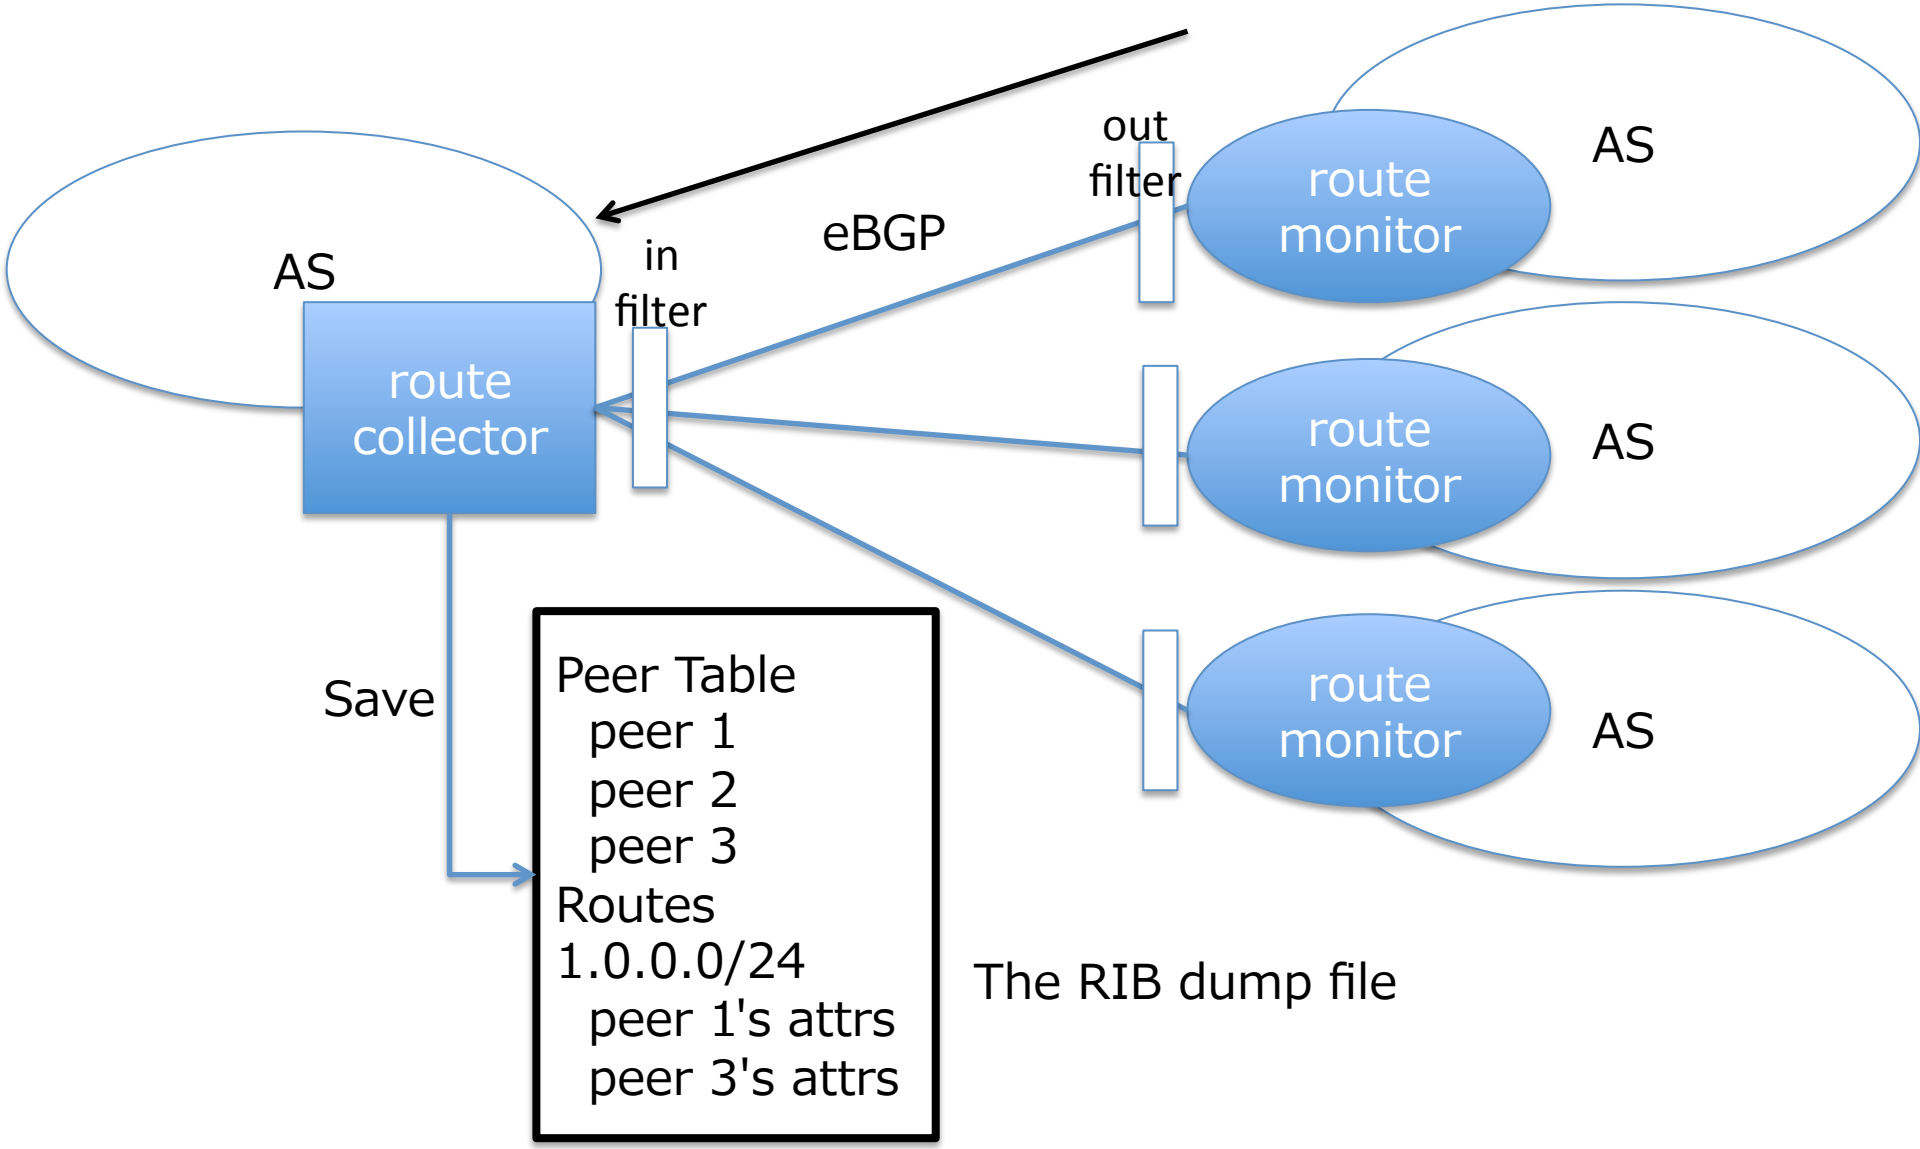

simple display

% ./src/bgpdump2 ../routeviews/oregon-ix2/rib.20150319.0000.bz2 | head -10

0.0.0.0/0 203.189.128.233 origin_as: 9902 as-path[2]: 23673 9902

1.0.0.0/24 213.144.128.203 origin_as: 15169 as-path[2]: 13030 15169

1.0.0.0/24 198.129.33.85 origin_as: 15169 as-path[2]: 293 15169

1.0.0.0/24 5.101.110.2 origin_as: 15169 as-path[3]: 202018 1299 15169

1.0.0.0/24 162.243.188.2 origin_as: 15169 as-path[2]: 393406 15169

1.0.0.0/24 95.85.0.2 origin_as: 15169 as-path[2]: 200130 15169

1.0.0.0/24 192.241.164.4 origin_as: 15169 as-path[2]: 62567 15169

1.0.0.0/24 129.250.0.11 origin_as: 15169 as-path[2]: 2914 15169

1.0.0.0/24 66.185.128.1 origin_as: 15169 as-path[2]: 1668 15169

1.0.0.0/24 173.205.57.234 origin_as: 15169 as-path[3]: 53364 3257 15169

Speed

74M routeviews/oregon-ix2/rib.20150319.0000.bz2

% /usr/bin/time -p bash -c 'bzcat routeviews/oregon-ix2/rib.20150319.0000.bz2 > /dev/null 2>&1 '

real 28.98 user 28.80 sys 0.10

% /usr/bin/time -p bash -c 'bzcat routeviews/oregon-ix2/rib.20150319.0000.bz2 | ./zebra-dump-parser/zebra-dump-parser.pl > /dev/null 2>&1 '

real 428.29 user 467.67 sys 1.07

% /usr/bin/time -p bash -c './libbgpdump-1.4.99.11/bgpdump routeviews/oregon-ix2/rib.20150319.0000.bz2 > /dev/null 2>&1 '

real 148.71 user 148.38 sys 0.24

% /usr/bin/time -p bash -c './bgpdump2/src/bgpdump2 routeviews/oregon-ix2/rib.20150319.0000.bz2 > /dev/null 2>&1 '

real 89.09 user 88.74 sys 0.19

display peer index table

% ./src/bgpdump2 ../routeviews/oregon-ix2/rib.20150319.0000.bz2 -P | head -12

Collector BGP ID: 128.223.51.102

View Name Length: 0

View Name:

Peer Count: 58

peer_table[0] changed: 0.0.0.0 asn:0 [129.250.0.11|::]

peer_table[1] changed: 10.10.10.252 asn:53364 [173.205.57.234|::]

peer_table[2] changed: 0.0.0.0 asn:0 [192.241.164.4|::]

peer_table[3] changed: 4.69.184.193 asn:3356 [4.69.184.193|::]

peer_table[4] changed: 5.101.110.2 asn:202018 [5.101.110.2|::]

peer_table[5] changed: 12.0.1.63 asn:7018 [12.0.1.63|::]

peer_table[6] changed: 64.57.28.241 asn:11537 [64.57.28.241|::]

peer_table[7] changed: 66.185.128.1 asn:1668 [66.185.128.1|::]

display per peer routing table

```
% ./src/bgpdump2 ../routeviews/oregon-ix2/rib.20150319.0000.bz2 -p 19 | head -8
1.0.0.0/24 129.250.0.11 origin_as: 15169 as-path[2]: 2914 15169
1.0.4.0/24 129.250.0.11 origin_as: 56203 as-path[4]: 2914 3257 4826 56203
1.0.5.0/24 129.250.0.11 origin_as: 56203 as-path[4]: 2914 3257 4826 56203
1.0.6.0/24 129.250.0.11 origin_as: 56203 as-path[4]: 2914 3257 4826 56203
1.0.7.0/24 129.250.0.11 origin_as: 56203 as-path[6]: 2914 3257 4826 56203 56203 56203
1.0.38.0/24 129.250.0.11 origin_as: 24155 as-path[2]: 2914 24155
1.0.43.0/24 129.250.0.11 origin_as: 24155 as-path[4]: 2914 1299 10026 24155
1.0.44.0/24 129.250.0.11 origin_as: 24155 as-path[4]: 2914 1299 10026 24155
```

per peer statistics

```
% ./src/bgpdump2 ../routeviews/oregon-ix2/rib.20150319.0000.bz2 -k | less
      :
peer[19]:
Number of routes: 524728
Number of routes per plen:
   /0 :      0    /1 :      0    /2 :      0    /3 :      0    /4 :      0
   /5 :      0    /6 :      0    /7 :      0    /8 :     16    /9 :     12
   /10:     33    /11:     92    /12:    263    /13:    501    /14:    991
   /15:   1701    /16:  12887    /17:   7047    /18:  11927    /19:  24655
   /20:  35315    /21:  37702    /22:  57644    /23:  49291    /24: 284473
   /25:     52    /26:     51    /27:     20    /28:     20    /29:     16
   /30:      6    /31:      0    /32:     13
Number of nexthops: 1
Number of origin_as: 49440
Number of unique as paths: 75351
Number of as path len: 39
      :
```

Distribution in Prefix Length

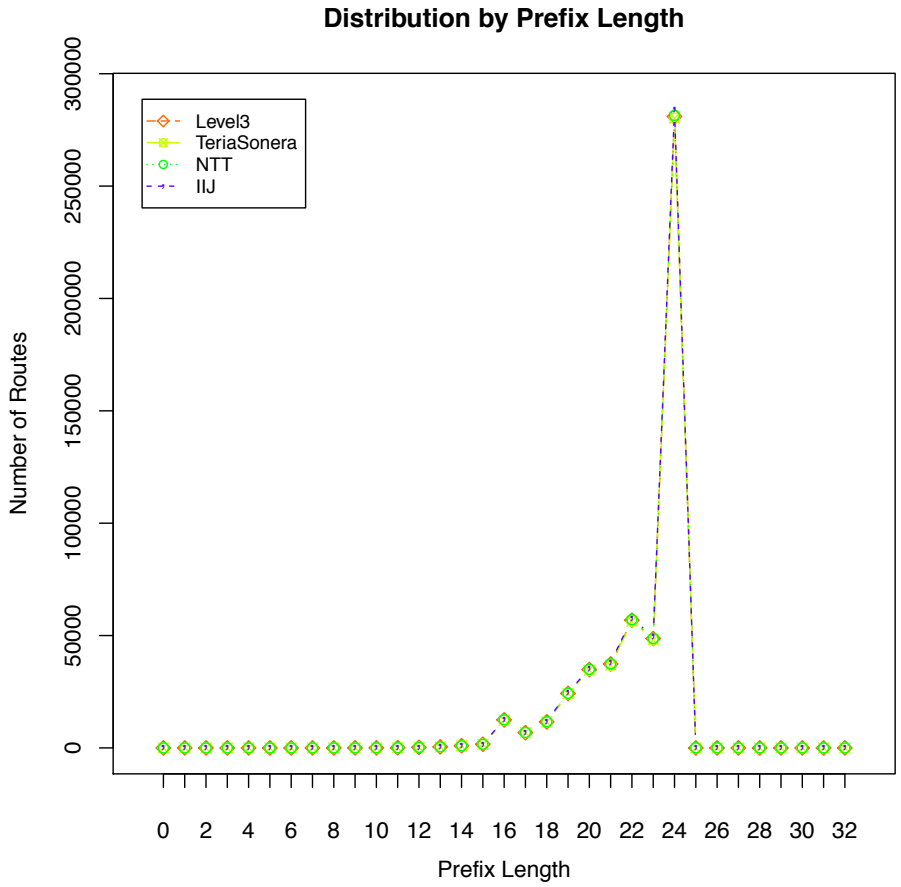

routing table lookup (longest-match)

- Routing Table (PATRICIA) can be created.
- one address lookup.

% ./src/bgpdump2 ../routeviews/oregon-ix2/rib.20150319.0000.bz2 -p 19 -l 8.8.8.8

looking up an address: 8.8.8.8

8.8.8.0/24 129.250.0.11 origin_as: 15169 as-path[2]: 2914 15169

- a list of address lookup contained in a file.

% ./src/bgpdump2 ../routeviews/oregon-ix2/rib.20150319.0000.bz2 -p 19 -L ../test-lookup.txt

50.0.0.0/16 129.250.0.11 origin_as: 7065 as-path[3]: 2914 8121 7065

100.0.0.0/16 129.250.0.11 origin_as: 701 as-path[2]: 2914 701

150.0.0.0: no route found.

200.0.0.0: no route found.

Diffs a peer against others

- script/compare-peer-with-others.sh
- marks

udiff:

<: the prefix is in left and unreachable in right (maybe partially).

+: the prefix is only in right (but it is reachable in left).

): the prefix is in right, and is covered by a shorter prefix in left that is unreachable in right. (i.e., the shorter is '<')

>: the prefix is in right and unreachable in left (maybe partially).

-: the prefix is only in left (but it is reachable in right).

(: the prefix is in left, and is covered by a shorter prefix in right that is unreachable in left. (i.e., the shorter is '>')

Comparison among peers

- diff routes between 2914 and 3356:

>70.61.0.0/20 4.69.184.193 origin_as: 10796 as-path[3]: 3356 7843 10796
>70.61.1.0/24 4.69.184.193 origin_as: 10796 as-path[3]: 3356 7843 10796
>70.61.2.0/24 4.69.184.193 origin_as: 30628 as-path[4]: 3356 7843 10796 30628
(70.61.4.0/24 129.250.0.11 origin_as: 10796 as-path[4]: 2914 2828 7843 10796
(70.61.5.0/24 129.250.0.11 origin_as: 10796 as-path[4]: 2914 2828 7843 10796
(70.61.6.0/24 129.250.0.11 origin_as: 33363 as-path[5]: 2914 2828 7843 33363 33363
+71.29.112.0/21 4.69.184.193 origin_as: 7029 as-path[3]: 3356 2828 7029
>71.44.17.0/24 4.69.184.193 origin_as: 33363 as-path[2]: 3356 33363
>71.44.53.0/24 4.69.184.193 origin_as: 33363 as-path[2]: 3356 33363
>71.44.62.0/24 4.69.184.193 origin_as: 33363 as-path[2]: 3356 33363
-71.252.67.0/24 129.250.0.11 origin_as: 64512 as-path[3]: 2914 701 64512
+72.0.224.0/23 4.69.184.193 origin_as: 19940 as-path[2]: 3356 19940
+72.0.227.0/24 4.69.184.193 origin_as: 19940 as-path[2]: 3356 19940

missing prefix ranking in AS 2914 in Oregon IX

- cat data/oregon-ix2-rib.20150319.0000.bz2-p19/route-diff-rib.20150319.0000.bz2-p19-*-diff.txt | grep '^>' | awk '{print $1;}' | sort -n | uniq -c | sort -n –r

38 >94.176.2.0/24 94.176.2.0 ASTIMP-ASAstimpConsultingSRL,RO
38 >94.176.131.0/24 94.176.131.0 VOXILITY-ASVoxilityS.R.L.,RO
38 >94.156.77.0/24 94.156.77.0 NETERRA-ASNeterraLtd.,BG
38 >94.156.185.0/24 94.156.185.0 NETERRA-ASNeterraLtd.,BG
38 >94.156.184.0/24 94.156.184.0 NETERRA-ASNeterraLtd.,BG
38 >93.123.18.0/24 93.123.18.0 NETERRA-ASNeterraLtd.,BG
38 >93.120.36.0/22 93.120.36.0 VOXILITY-ASVoxilityS.R.L.,RO
38 >93.120.35.0/24 93.120.35.0 ASTIMP-ASAstimpConsultingSRL,RO
38 >93.115.92.0/22 93.115.92.0 VOXILITY-ASVoxilityS.R.L.,RO
38 >93.115.88.0/22 93.115.88.0 VOXILITY-ASVoxilityS.R.L.,RO

ranking in appearance count of whois descr

180
88 TFN-TWTaiwanFixedNetwork,TelcoandNetworkServiceProvider.,TW
60 VOXILITY-ASVoxilityS.R.L.,RO
60 FPT-AS-APTheCorporationforFinancing&PromotingTechnology,VN
42 SCRR-10796-TimeWarnerCableInternetLLC,US
37 TAIWANMOBILE-ASTaiwanMobileCo.,Ltd.,TW
33 XTGLOBALXTGLOBALNETWORKSLTD.,RO
32 DATAFRAMELO-DataframeLogistics,Inc.,US
28 ONE-NET-HKINTERNET-SOLUTION-HK,CN
24 VNPT-AS-VNVNPTCorp,VN

2914:4429 community tag: do not advertise in Asia.
http://www.us.ntt.net/support/policy/routing.cfm

full route number

% sh full-route-number.sh -p 0 -p 1 -p 2 -p 3 -p 5 -p 7 -p 11 –p 13
../../../routeviews/oregon2-summary/rib.20140{1,2,3}*.bz2

- peer 13 was empty.
- The R script is also provided.

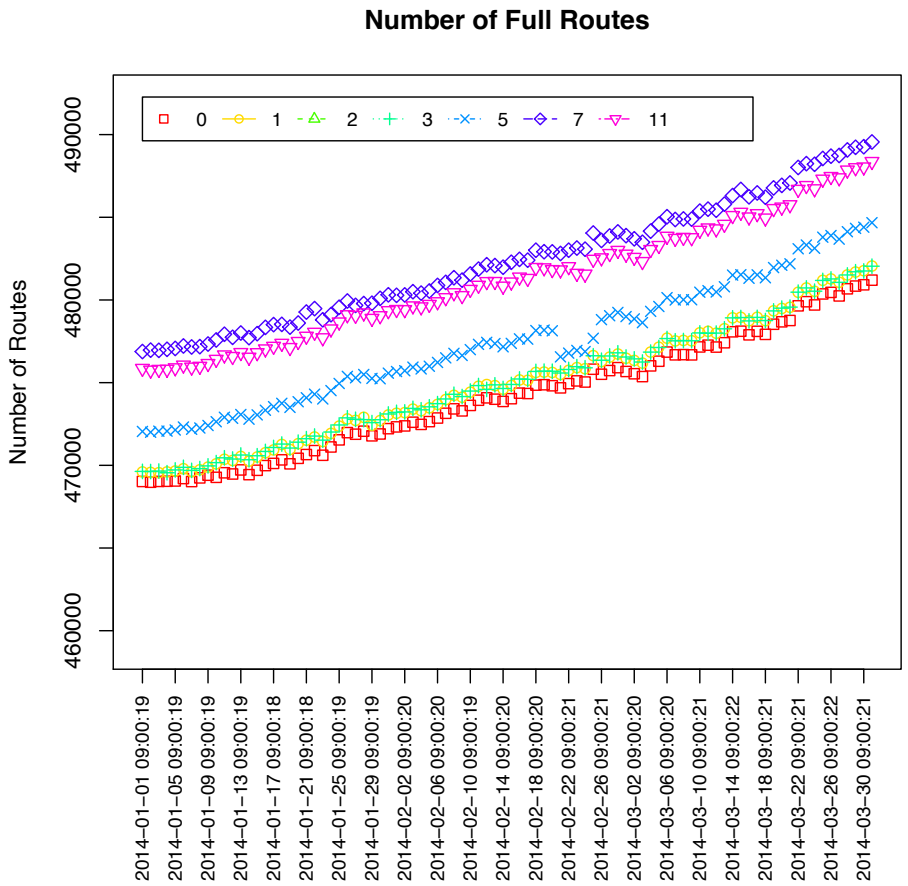

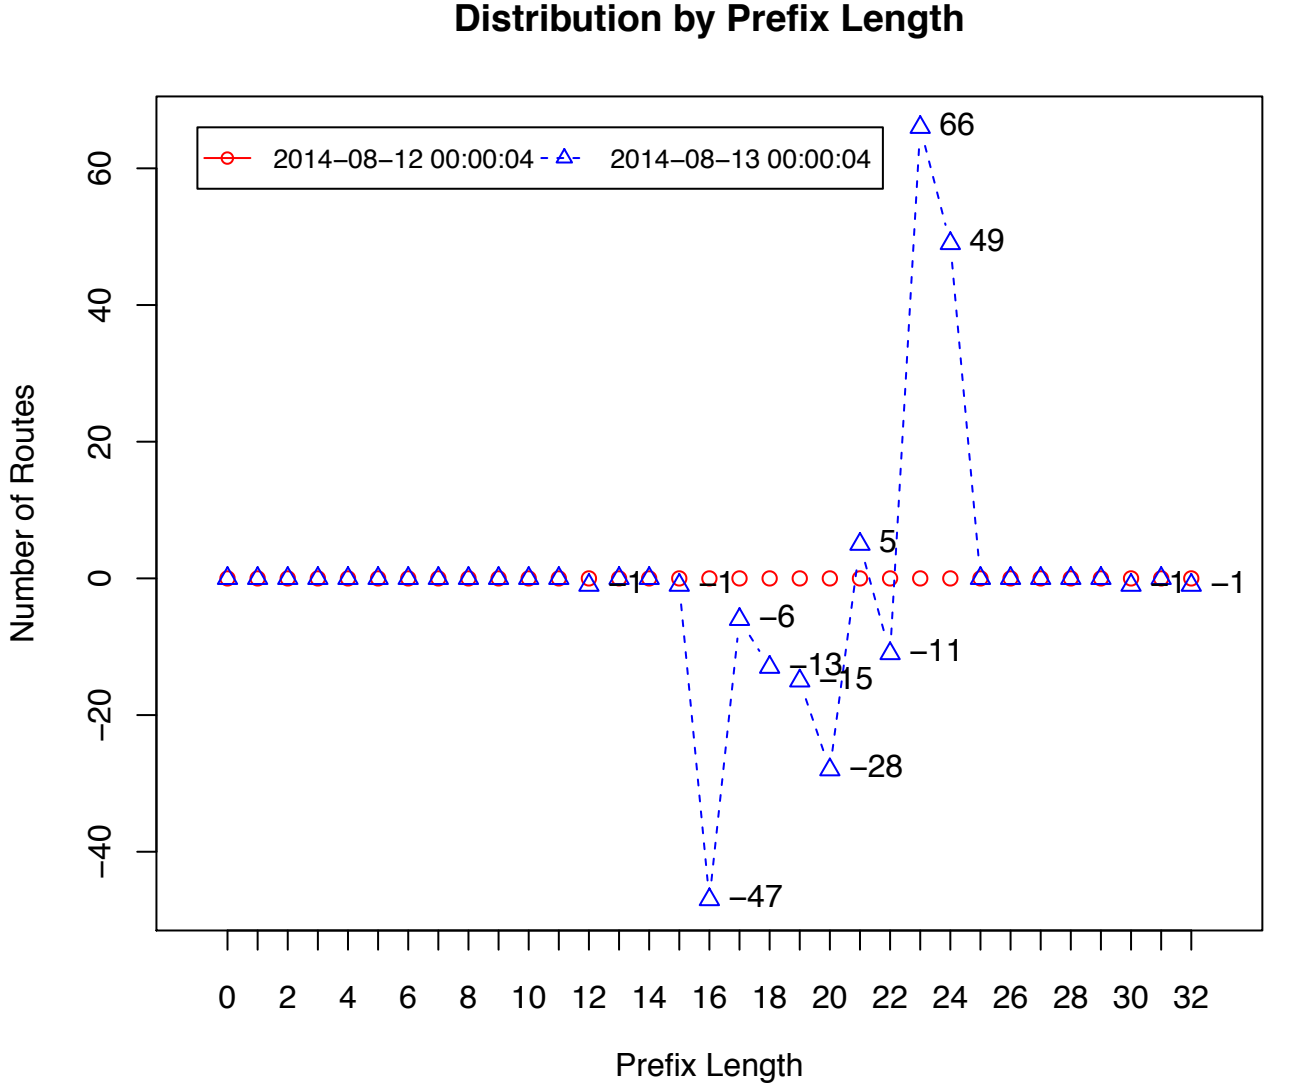

Increase in the number of route.
Disbribution per prefix length.
2014-08-12 00:00 – 2014-08-13 00:00

A Use Case (1)

- Access source analysis of an open NTP server
  - tcpdump: 50M pcap packets
  - Solving 50M source IP addresses to Origin ASes
  - Whois/DNS query to Team Cymru didn't work.
  - bgpdump2 helped it:
  - The user's voice: It was good because
    - the portability (local resolution from the file),
    - the speed was super quick,
    - will help in retrospect activities.

A Use Case (2)

- 50 files, each including approx. 1M addresses
  - time bgpdump2 $HOME/work/bgp/rib/rib.20150508.0600 -p 0 -L <filename>
- using raw RIB file, cached on memory
- Took 00:03:59.1 (in total) to solve 49,921,136 IP addresses
  - 3.912u 0.848s 0:05.76 82.4%0+0k 128+28io 8pf+0w
  - 3.908u 0.651s 0:04.59 99.1%0+0k 38+6io 0pf+0w
  - 4.063u 0.689s 0:04.78 99.1%0+0k 112+13io 0pf+0w
  - 3.969u 0.646s 0:04.64 99.1%0+0k 4+6io 0pf+0w
  - 4.002u 0.650s 0:04.70 98.9%0+0k 8+4io 0pf+0w
  - 4.038u 0.646s 0:04.72 98.9%0+0k 0+4io 0pf+0w
  - 4.136u 0.661s 0:04.81 99.5%0+0k 2+17io 0pf+0w
  - :

Heatmap

- Just represents the density of IP reachability
  - A.B/16 depicts one point
  - In Hilbert Curve
  - lookup A.B.C.0 (head address of the /24) in the A.B/16 and counted the number of success (found)
  - A (the first octet) is labeled in white

20150319
p19 NTT

heatmap/oregon-ix2-rib.20150319.0000-p19

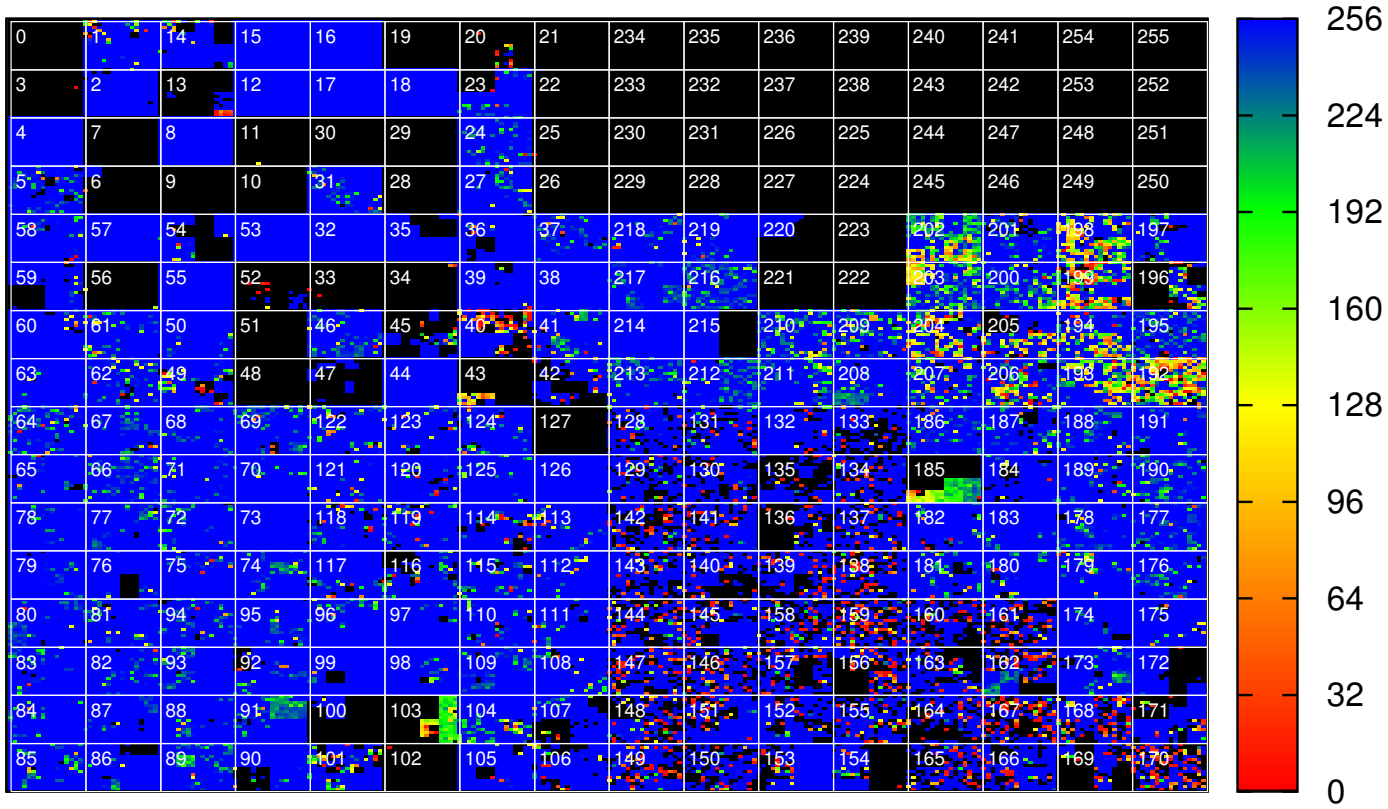

20150319
p3 Level3

heatmap/oregon-ix2-rib.20150319.0000-p3

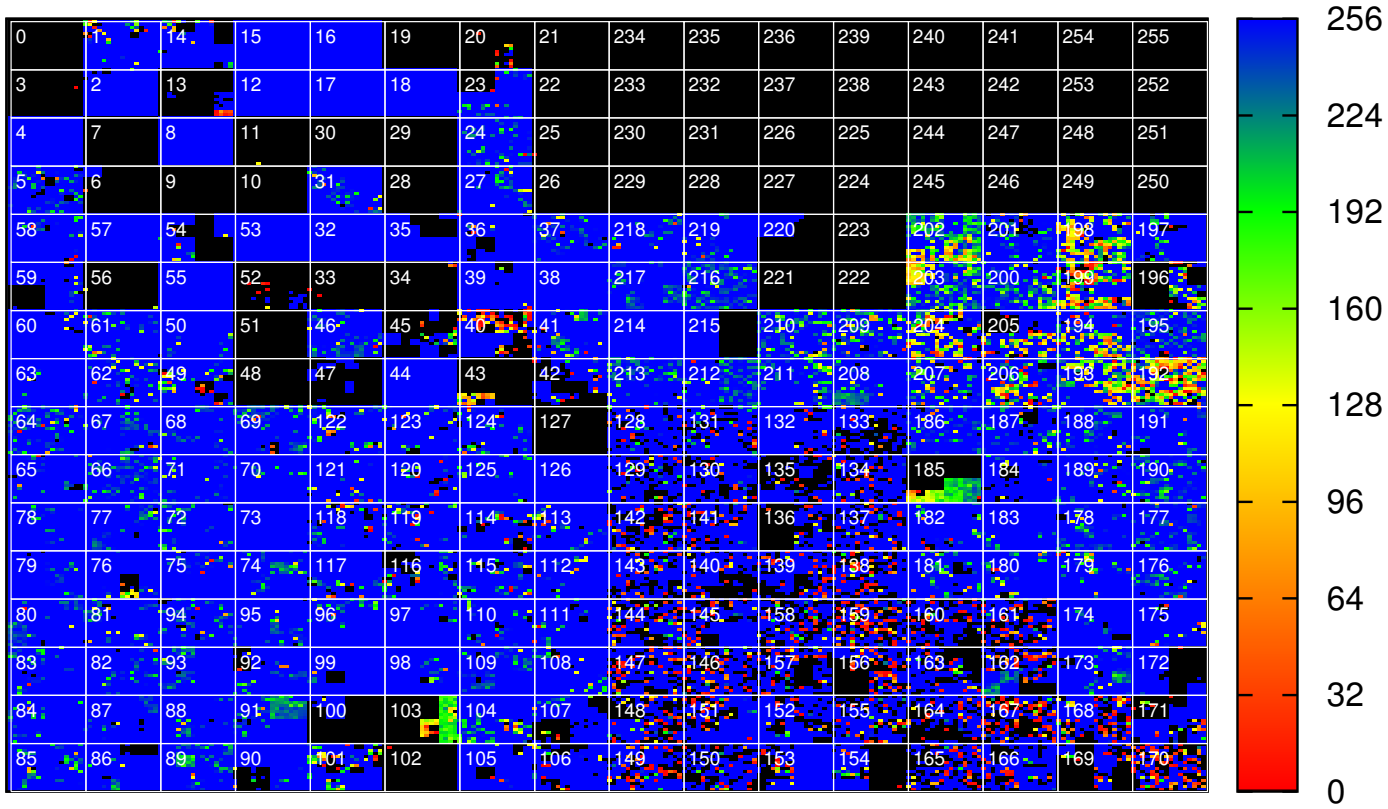

20150319
p1 ZEROFAIL

heatmap/oregon-ix2-rib.20150319.0000-p1

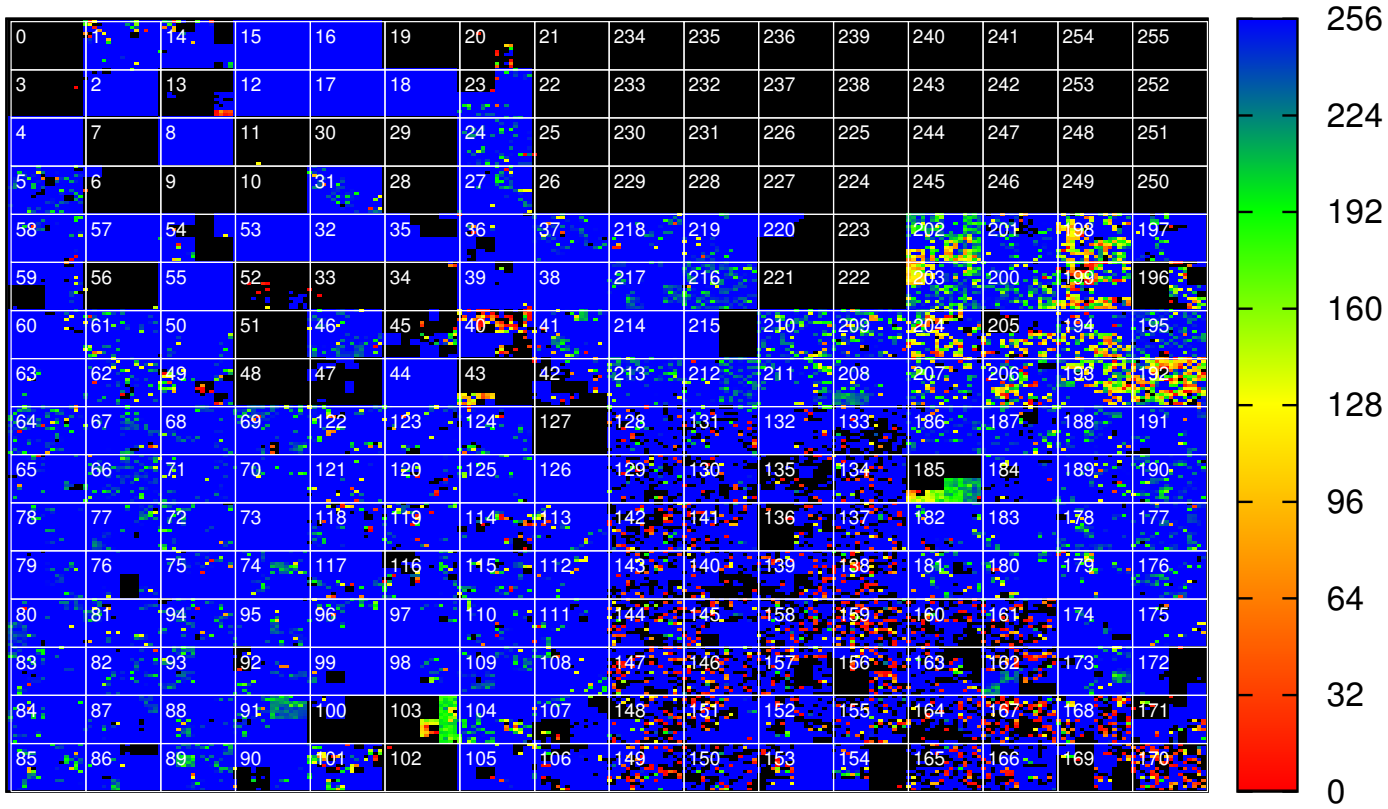

20150319
p4 Digital Ocean

heatmap/oregon-ix2-rib.20150319.0000-p4

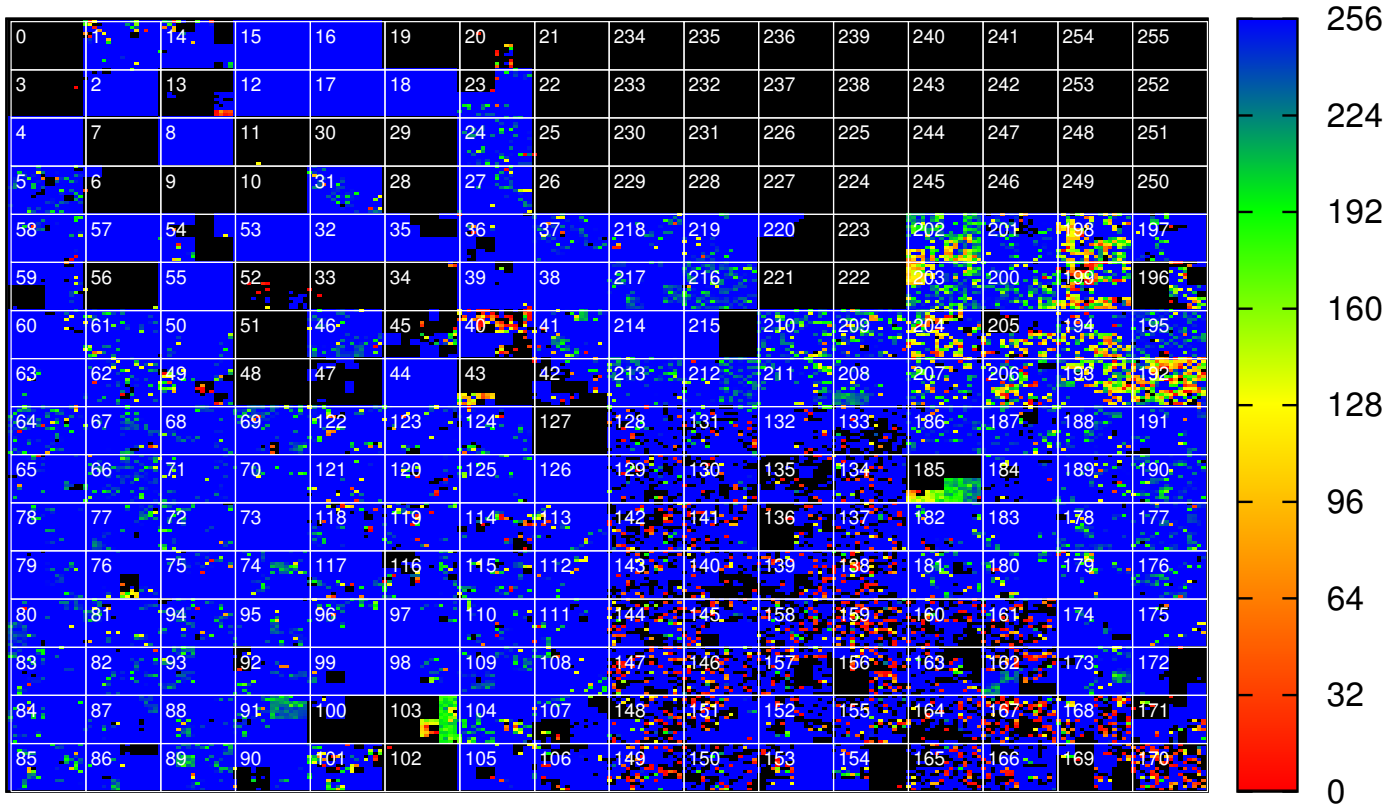

20150319
p5 AT&T

heatmap/oregon-ix2-rib.20150319.0000-p5

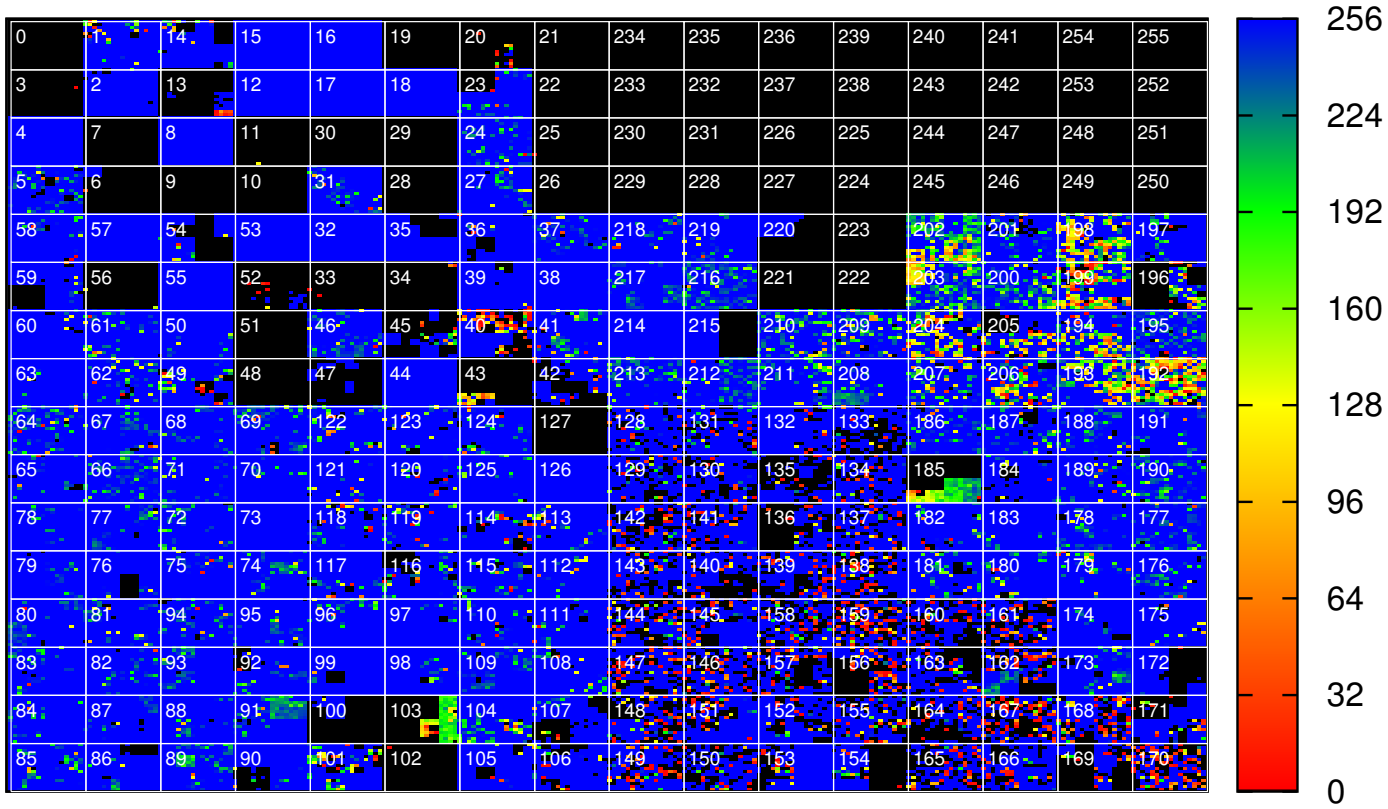

20150319
p7 AOL

heatmap/oregon-ix2-rib.20150319.0000-p7

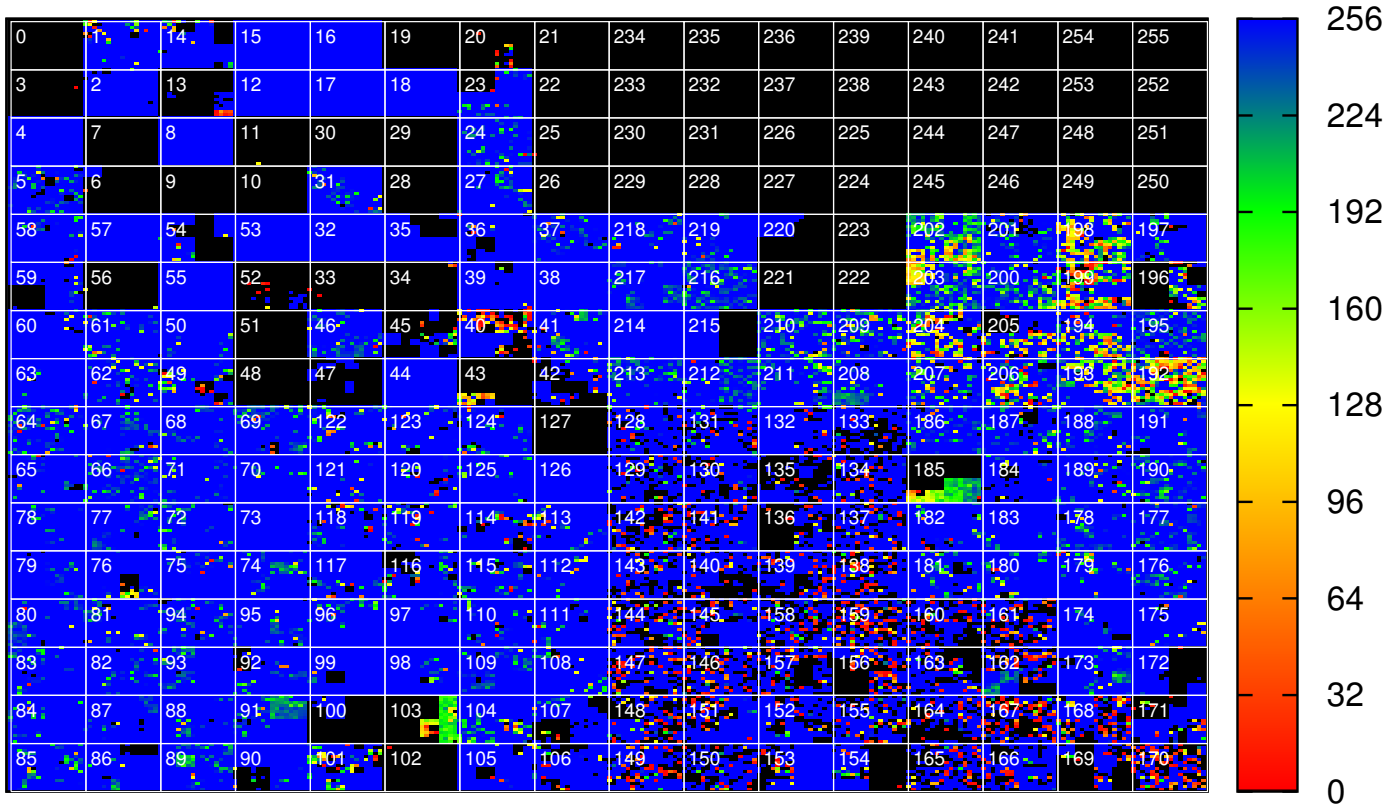

20150319
p8 Level3 (GBLX)

heatmap/oregon-ix2-rib.20150319.0000-p8

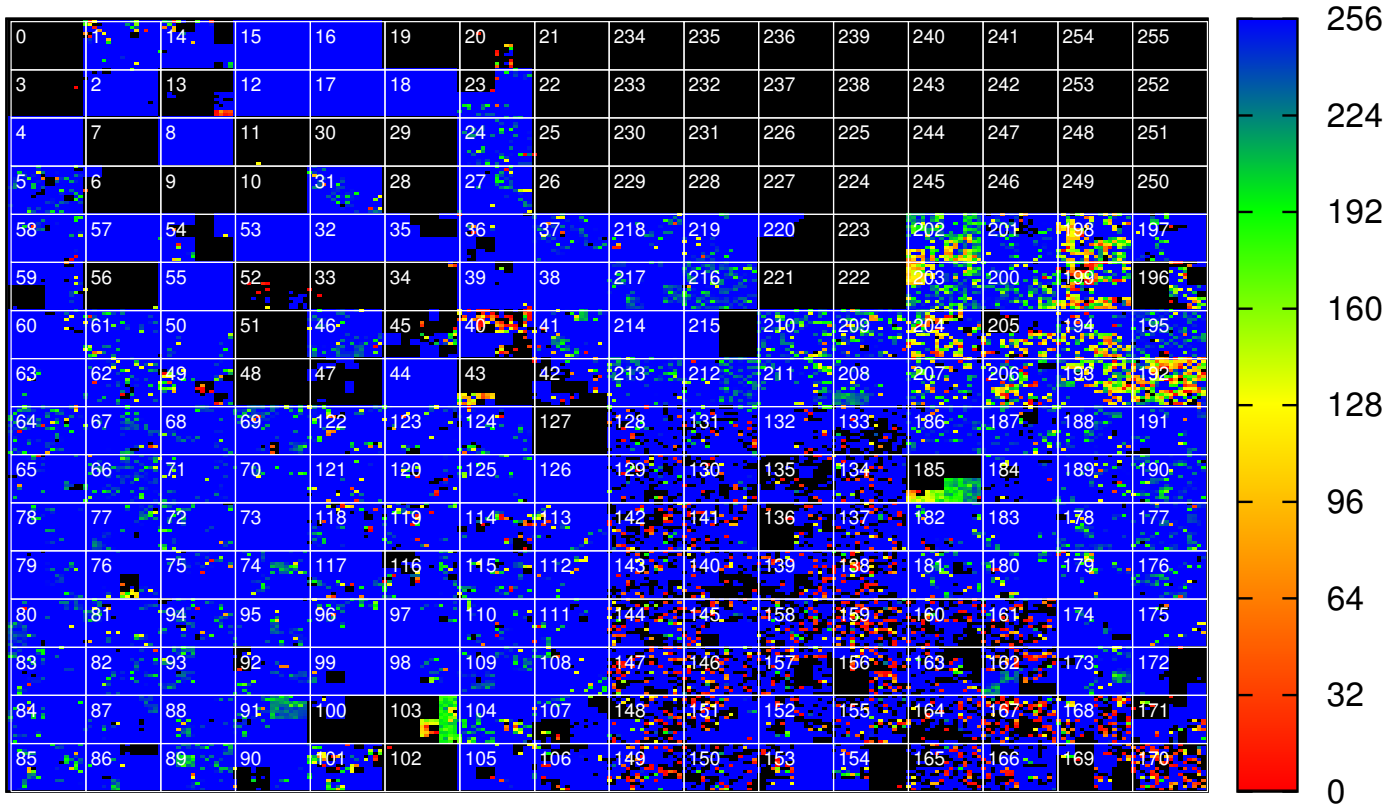

20150319
p9 Fibrenoire Internet

heatmap/oregon-ix2-rib.20150319.0000-p9

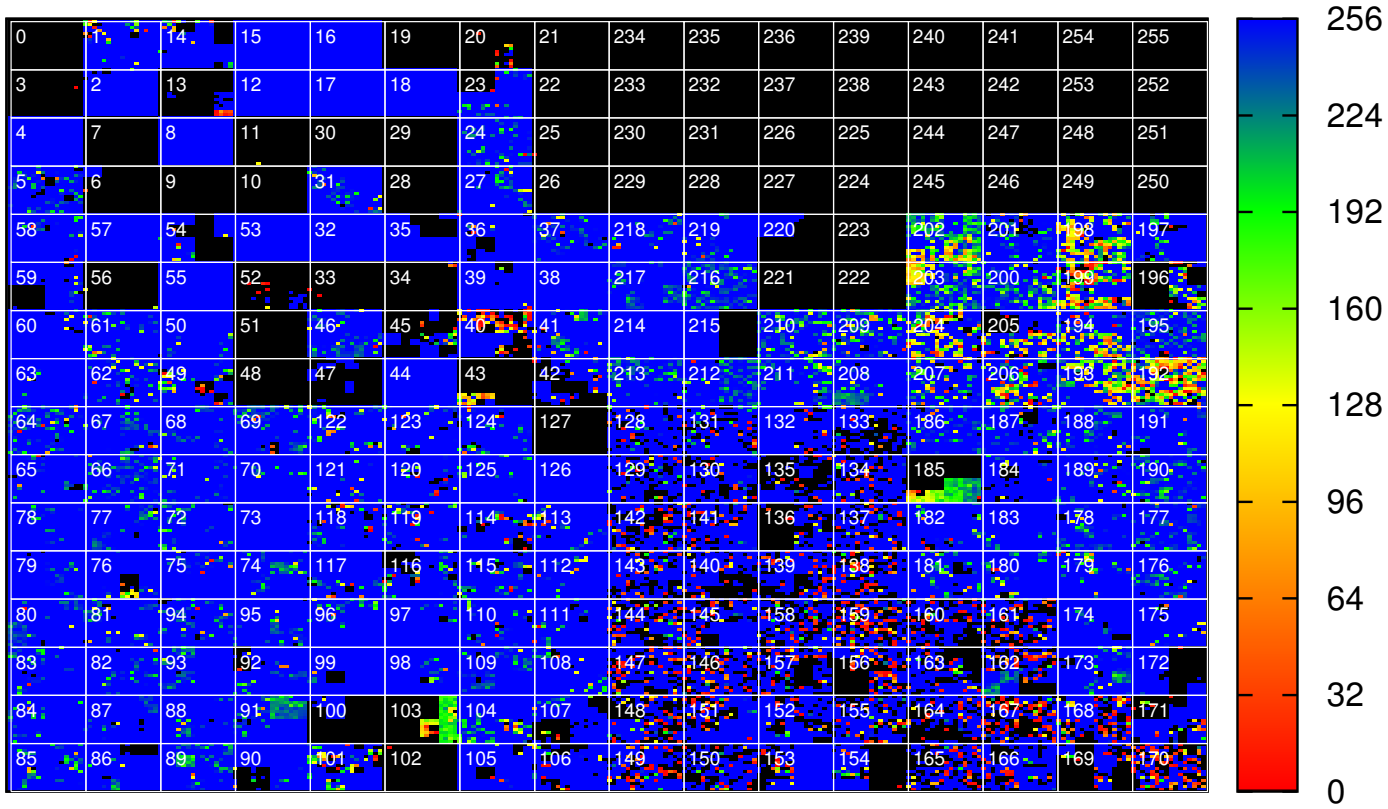

20150319
p11 TeriaSonera

heatmap/oregon-ix2-rib.20150319.0000-p11

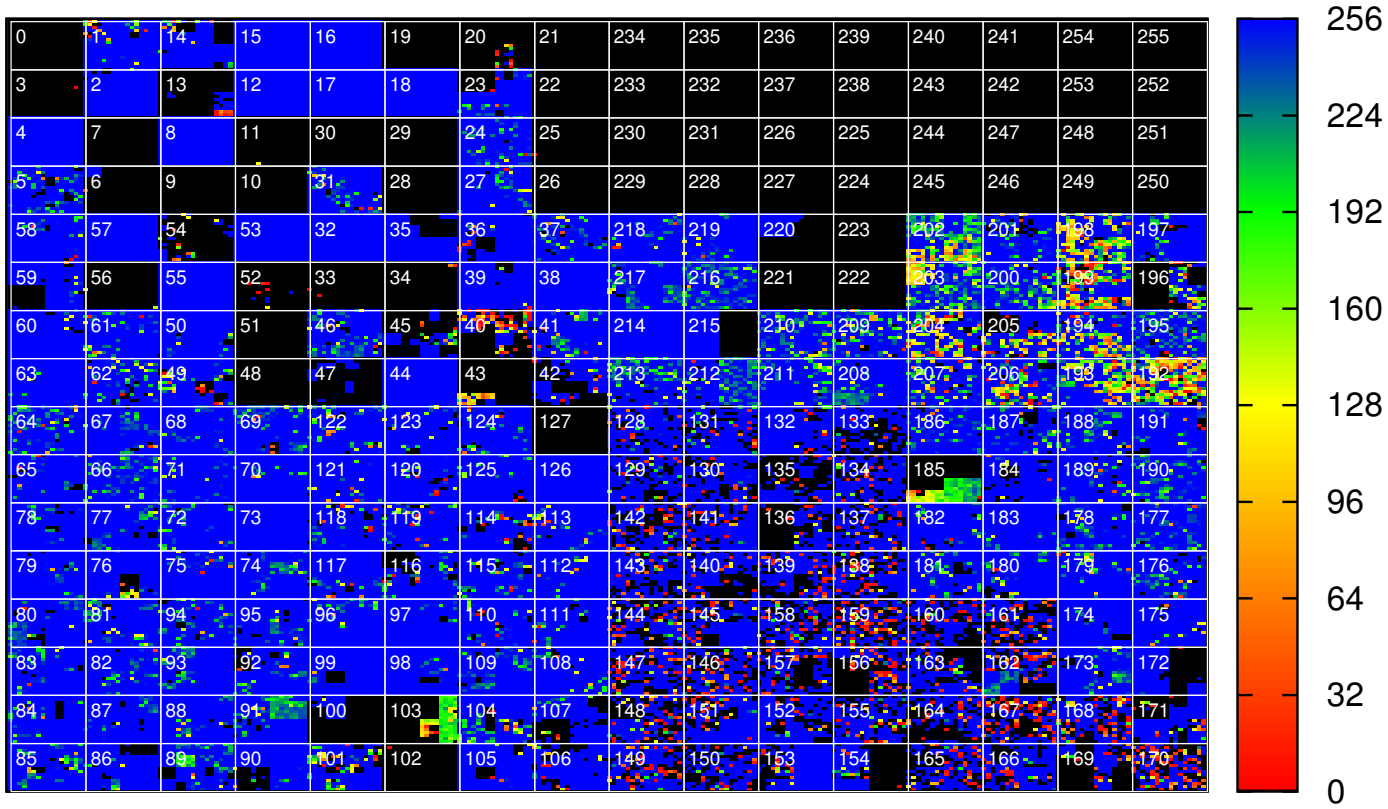

20150319
p12 OBIT

heatmap/oregon-ix2-rib.20150319.0000-p12

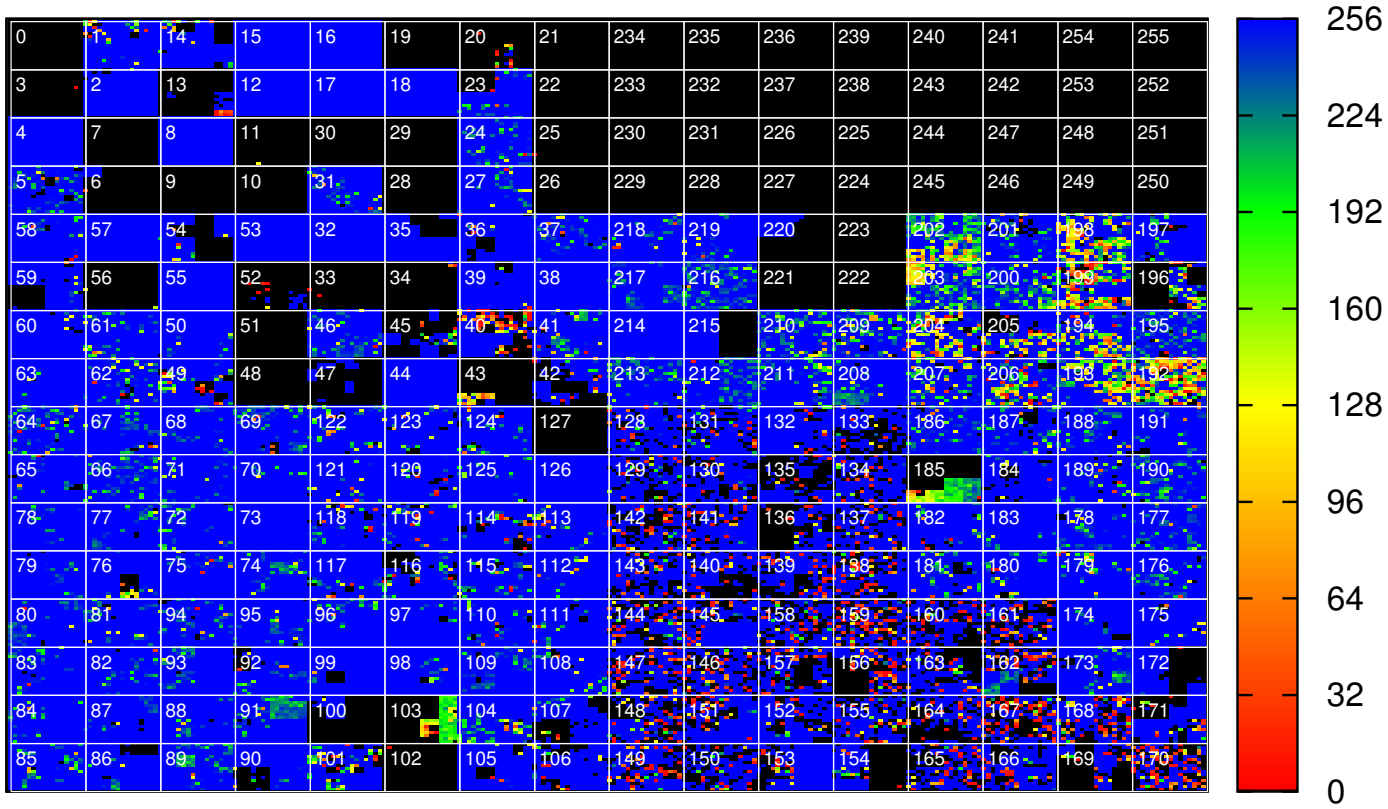

20150319
p42 IIJ

heatmap/oregon-ix2-rib.20150319.0000-p42

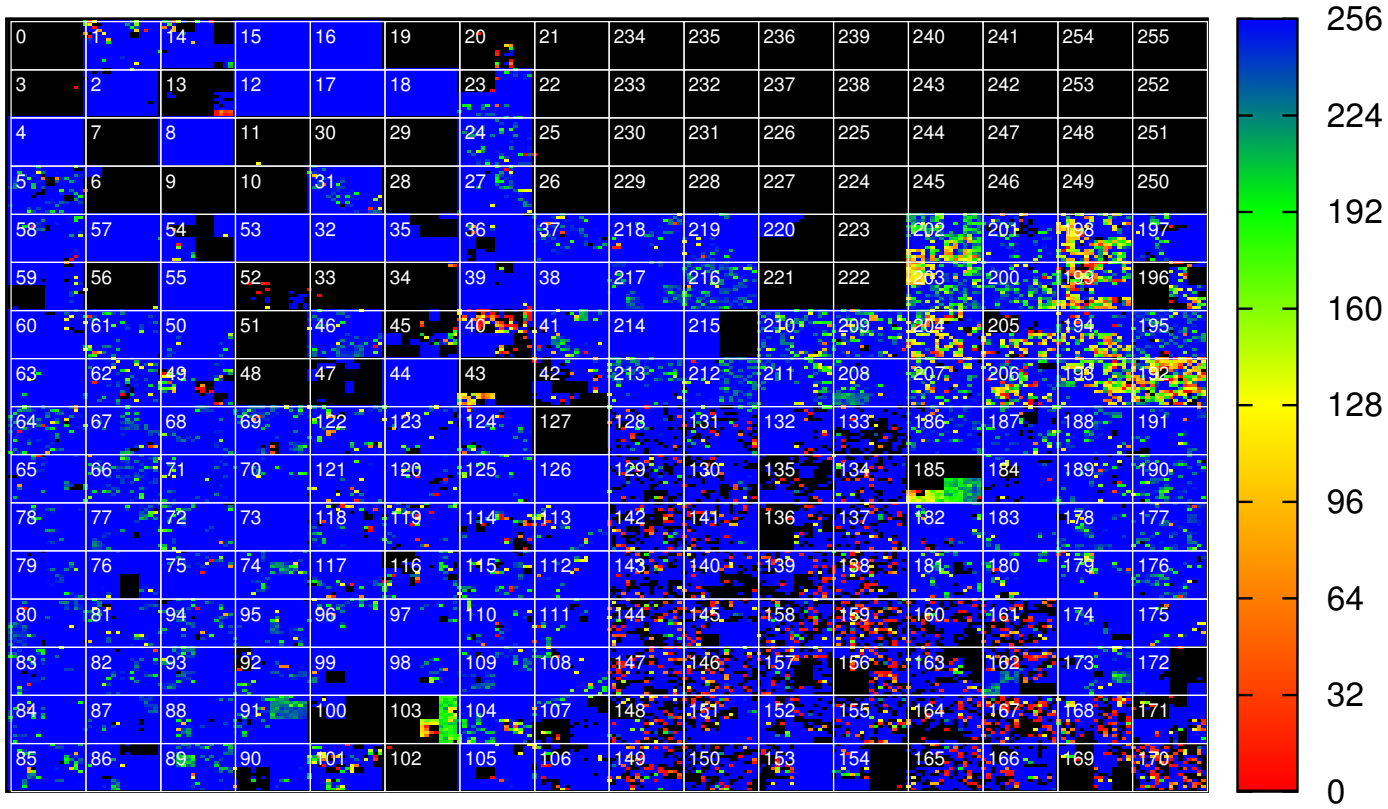

Timemachine: 2015 NTT

heatmap/oregon-ix2-rib.20150319.0000-p19

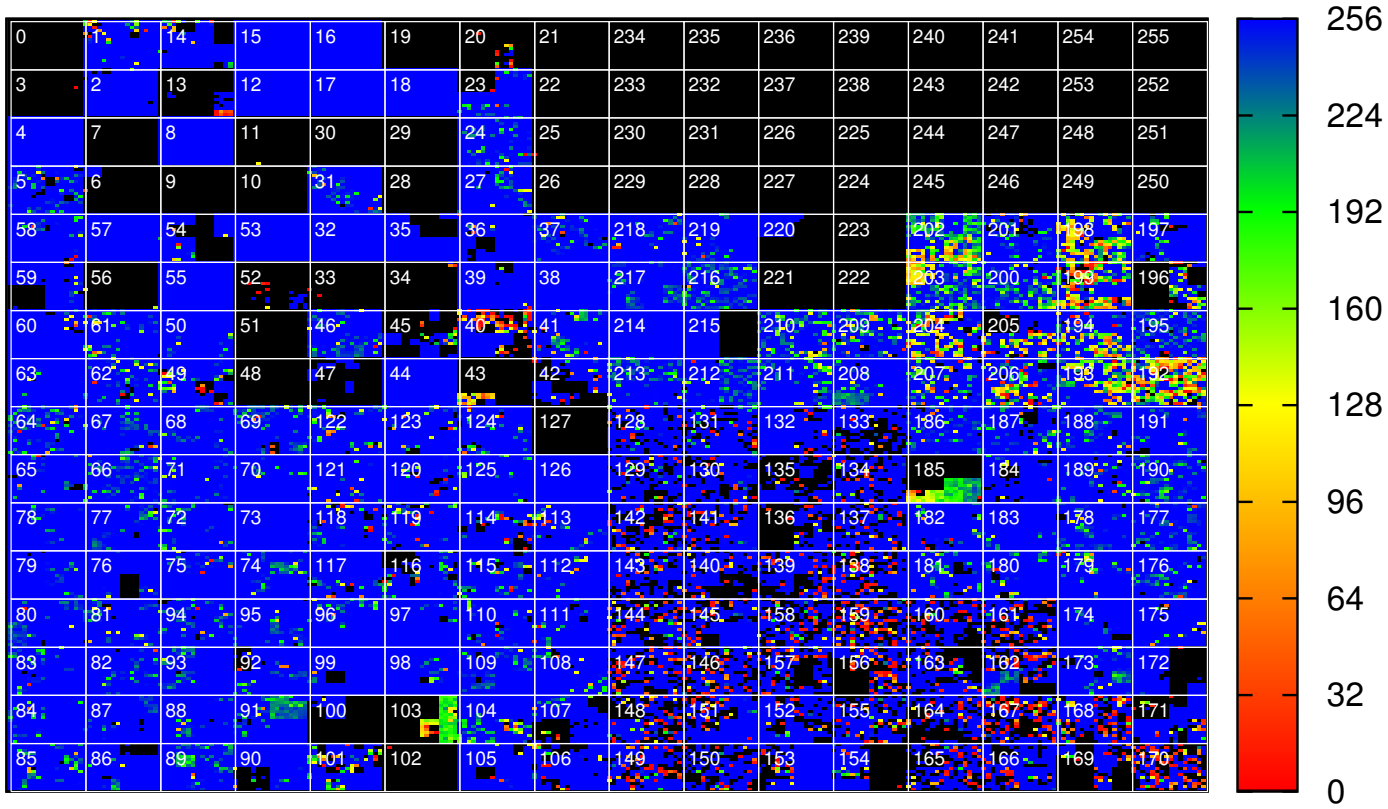

Timemachine: 2014 NTT

heatmap/oregon-ix2-rib.20140319.0000-p12

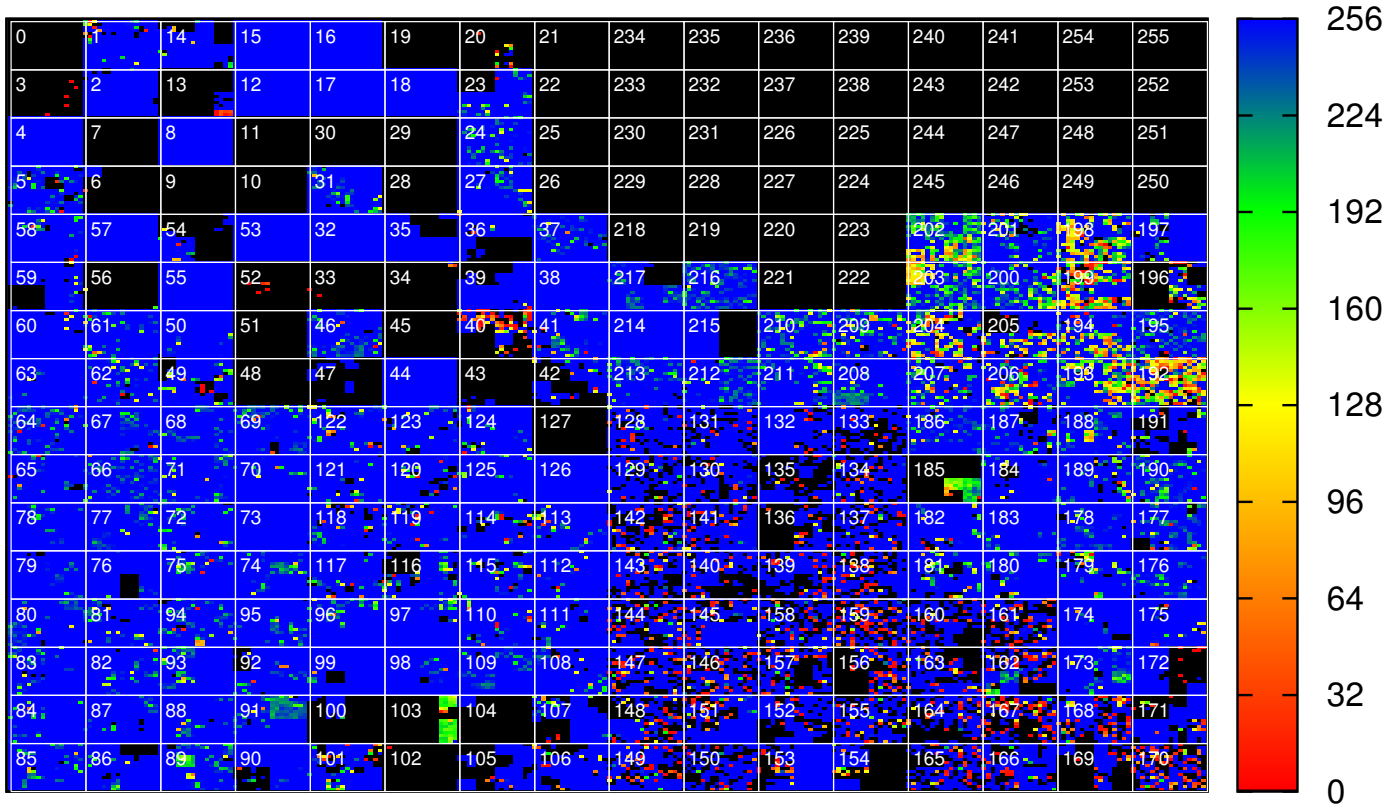

Timemachine: 2013 NTT

heatmap/oregon-ix2-rib.20130319.0000-p11

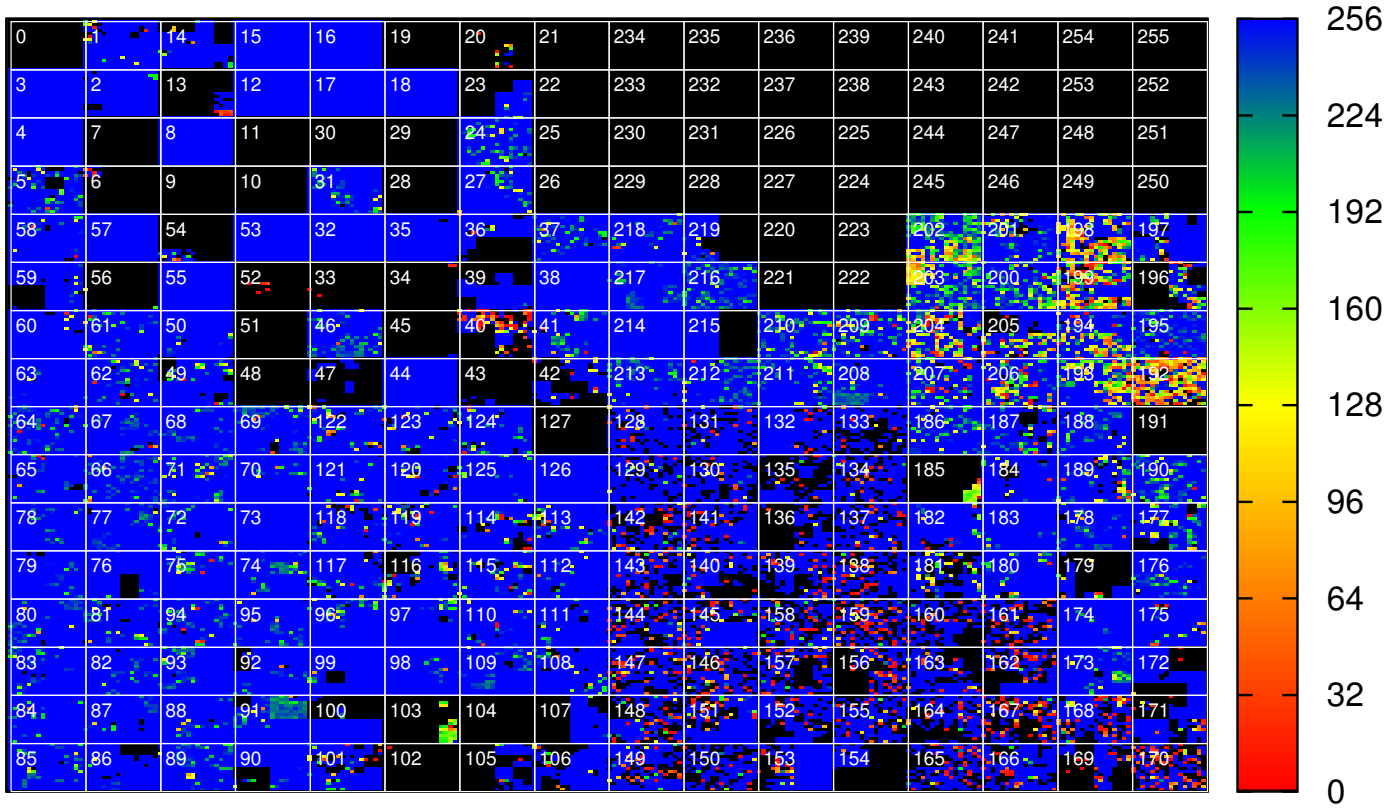

Timemachine: 2012 NTT

heatmap/oregon-ix2-rib.20120319.0000-p12

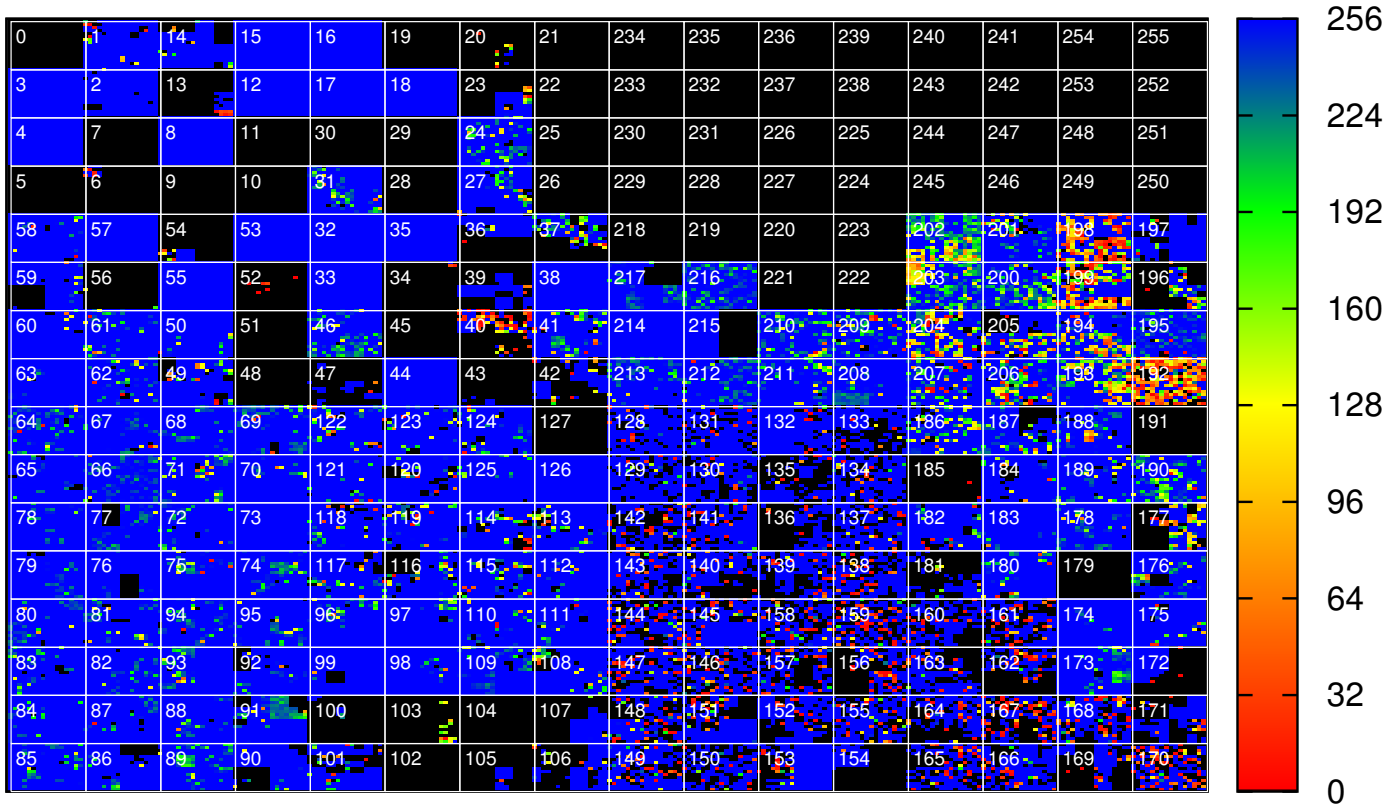

Timemachine: 2011 NTT

heatmap/oregon-ix2-rib.20110319.0000-p12

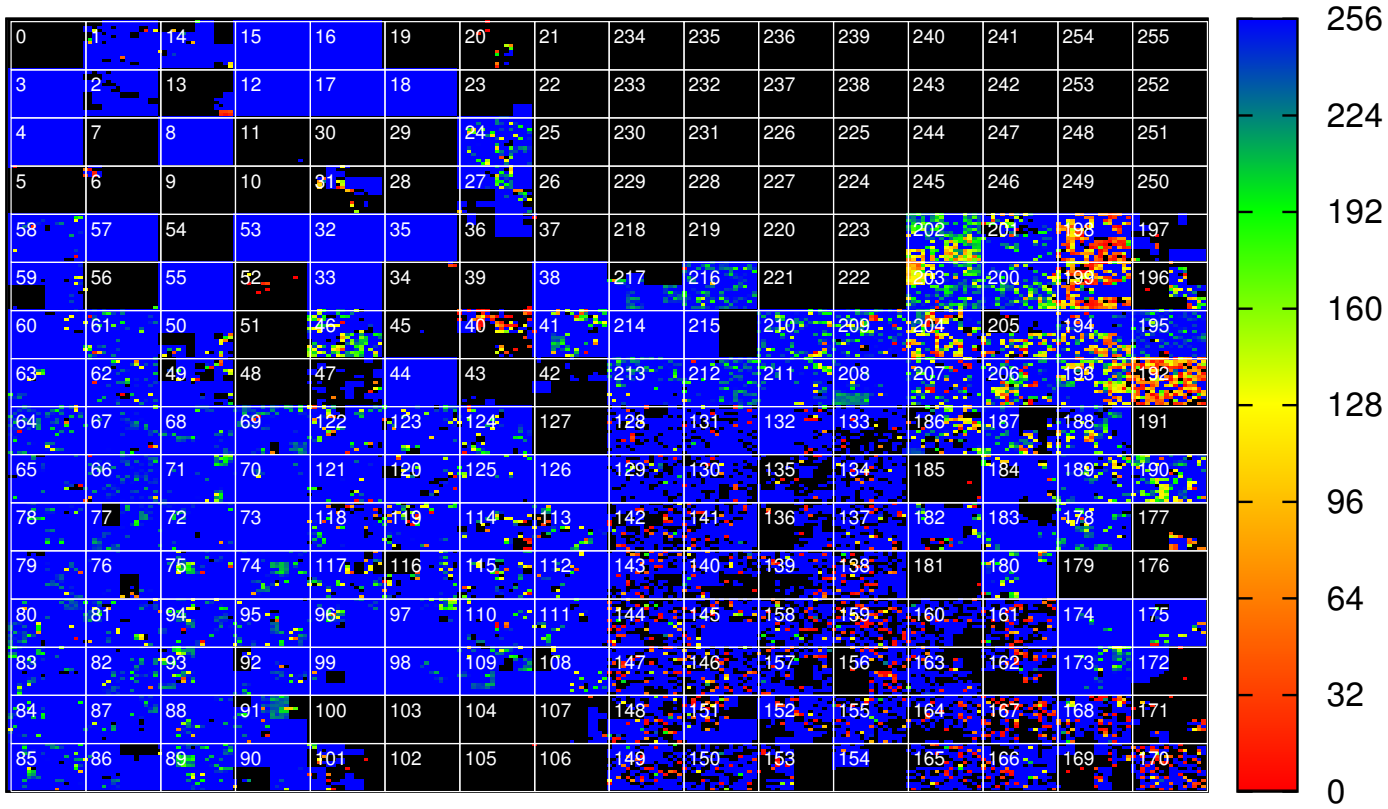

Timemachine: 2010 NTT

heatmap/oregon-ix2-rib.20100319.0000-p12

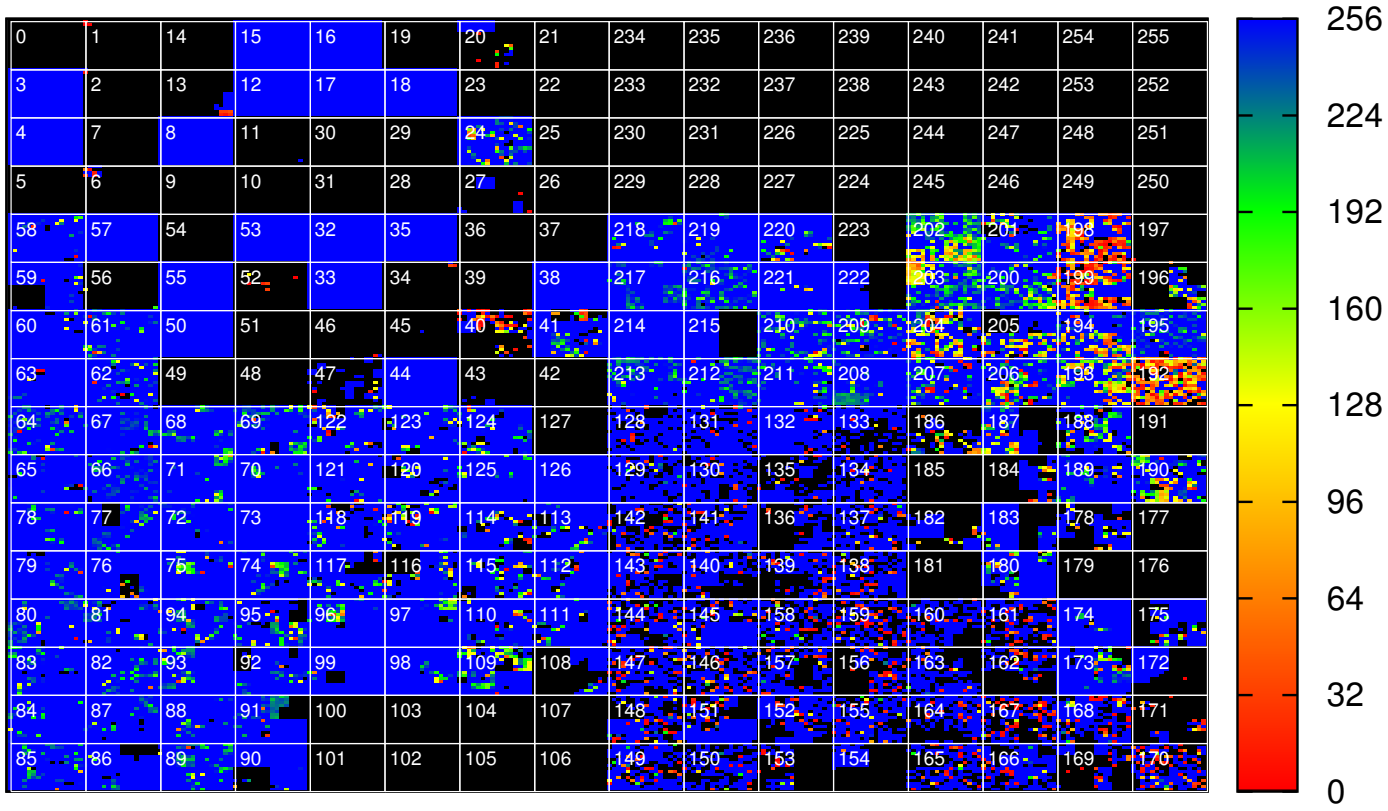

Wrap up

- An open-source tool is newly developed to analyze BGP dump: bgpdump2.
  - currently support only RIB dump (not BGP updates)
- New features
  - per peer statistics, per peer display of routes
  - diffs
  - PATRICIA-based routing lookup (longest-match)
- The tool should help:
  - BGP route debugging incl. route leaks
  - ISP comparison
  - analysis, and/or research

Thanks.
Questions ?

Existing tools

- libbgpdump
  - written in C. <https://bitbucket.org/ripencc/bgpdump/wiki/Home>
- zebra-dump-parser
  - written in Perl. <https://github.com/rfc1036/zebra-dump-parser>
- java-mrt library
  - written in Java. <https://github.com/paaguti/java-mrt>
- UCLA bgpparser
  - written in C++. <http://irl.cs.ucla.edu/software/bgpparser.html>
- mrtparse
  - written in Python. <https://github.com/YoshiyukiYamauchi/mrtparse>
- openbgpd bgpctl
  - written in C. <http://cvsweb.openbsd.org/cgi-bin/cvsweb/src/usr.sbin/bgpctl/>
- pybgpdump
  - writen in Python. <https://jon.oberheide.org/pybgpdump/>

libbgpdump compatible mode

- not finished: e.g., support for community, localpref, etc., are not yet completed.

```
% ./src/bgpdump2 ../routeviews/oregon-ix2/rib.20150319.0000.bz2 -m | head -3
TABLE_DUMP2|1426723200|B|203.189.128.233|23673|0.0.0.0/0|23673 9902|INCOMPLETE|203.189.128.233|0|0|0|NAG||
TABLE_DUMP2|1426723200|B|213.144.128.203|13030|1.0.0.0/24|13030 15169|INCOMPLETE|213.144.128.203|0|1|0|NAG||
TABLE_DUMP2|1426723200|B|198.129.33.85|293|1.0.0.0/24|293 15169|INCOMPLETE|198.129.33.85|0|0|0|NAG||
```