New BGP analysis tools and a look at the AS9121 Incident

Larry J. Blunk
Merit Network, Inc.

IEPG Meeting – 62nd IETF
Minneapolis
March 6, 2005

New BGP analysis tools

Merit is working to develop new tools for analysis of archived MRT data (such as from Routeviews and RIS)

Using libbgpdump for initial processing and analysis

- Reviewed libbgpdump code and made several fixes and performance improvements

Examining mechanisms for efficient aggregation and archival of BGP Update data

Using custom databases for optimized performance

Tools targeted at both researchers and for practical application by network operators

- Uses include examining hijackings, MOAS, flapping, martian/bogon announcements, etc.
- Also reachability issues and outages

Examining integration with Routing Registries for consistency checking and anomaly notification

Analyzing MRT Data

The Problem:

-Large volume of data

-Lots of data, little information (what does it all mean?)

-Lack of easy to use processing tools (only useful to researchers?)

Our Approach--BGP::Inspect

-Build a generic tool to preprocess MRT data and make it easier to query by everyone.

-Implement common queries to be fast, but also allow detailed data analysis if requested.

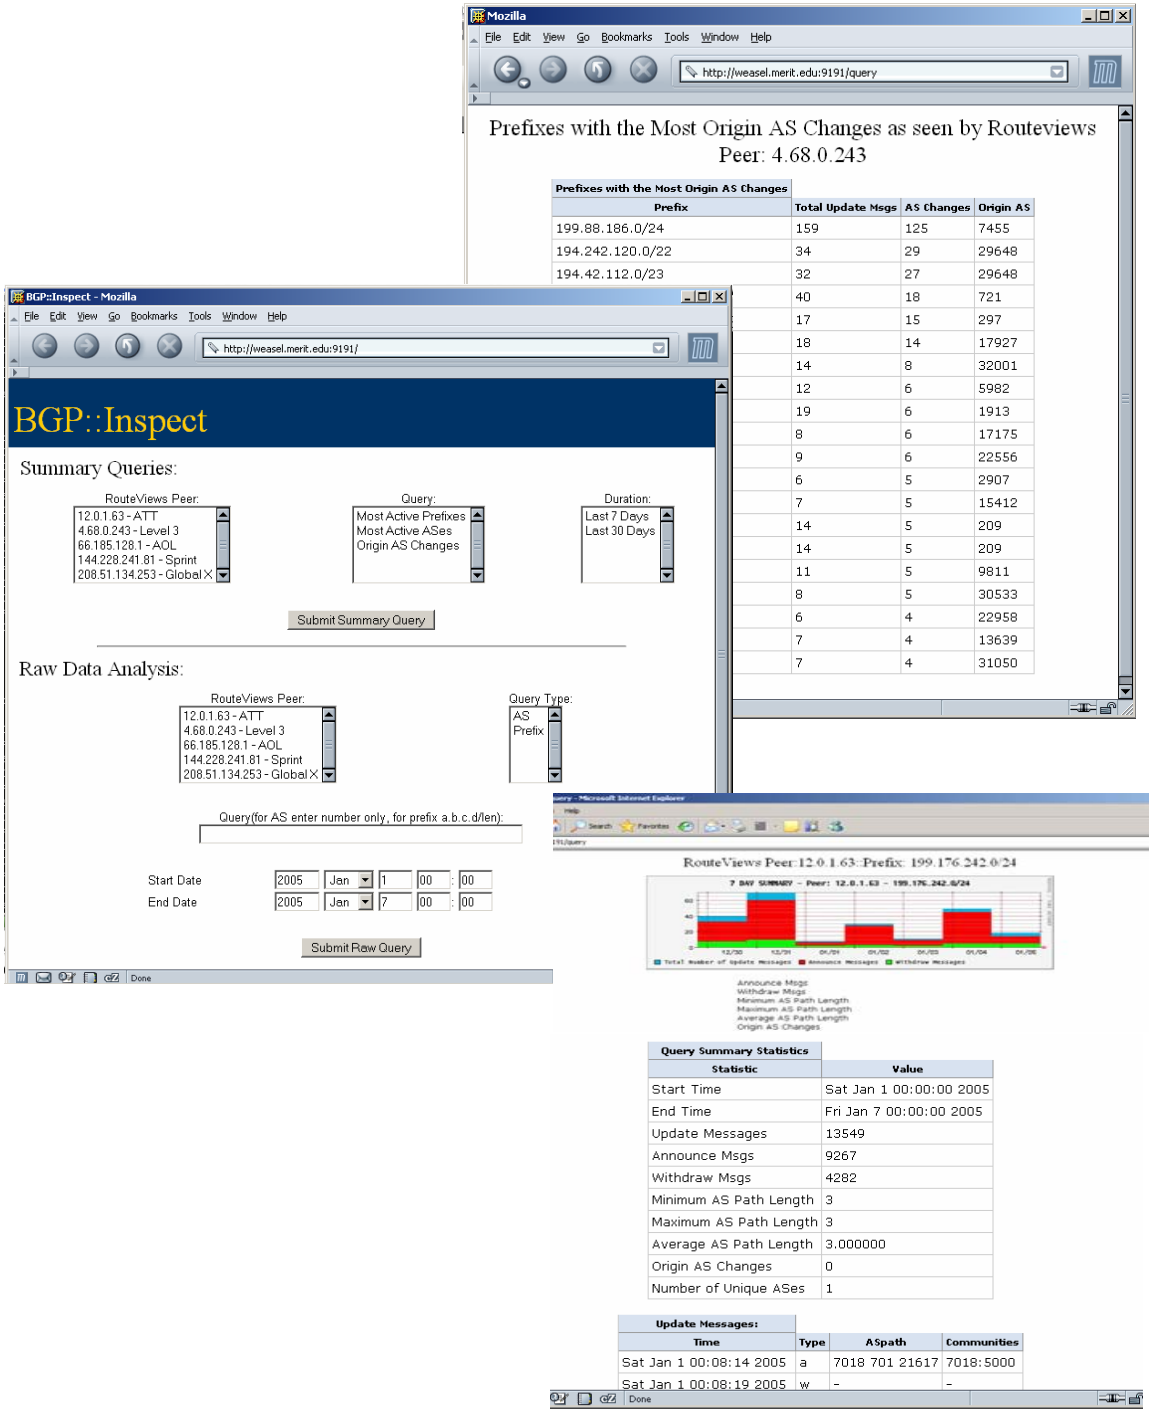

Prefixes with the Most Origin AS Changes			
Prefix	Total Update Msgs	AS Changes	Origin AS
199.88.186.0/24	159	125	7455
194.242.120.0/22	34	29	29648
194.42.112.0/23	32	27	29648
	40	18	721
	17	15	297
	18	14	17927
	14	8	32001
	12	6	5982
	19	6	1913
	8	6	17175
	9	6	22556
	6	5	2907
	7	5	15412
	14	5	209
	14	5	209
	11	5	9811
	8	5	30533
	6	4	22958
	7	4	13639
	7	4	31050

Query Summary Statistics	
Statistic	Value
Start Time	Sat Jan 1 00:00:00 2005
End Time	Fri Jan 7 00:00:00 2005
Update Messages	13549
Announce Msgs	9267
Withdraw Msgs	4282
Minimum AS Path Length	3
Maximum AS Path Length	3
Average AS Path Length	3.000000
Origin AS Changes	0
Number of Unique ASes	1

Update Messages:			
Time	Type	ASpath	Communities
Sat Jan 1 00:08:14 2005	a	7018 701 21617	7018:5000
Sat Jan 1 00:08:19 2005	w	-	-

BGP::Inspect

Key Ideas:

- Pre-process MRT data into easily query-able form
- Eliminate redundant data
- Use compression as necessary
- Pre-compute and store commonly useful statistics at data load time not at query-time

Current Status:

- Beta release of the tool at the end of January, limited data set, clean user query interface, moderately scalable, lots of interest from the networking community
- Next release scheduled for end of March, will include a more robust query front-end, a more scalable backend to allow large amounts of data to be pre-loaded, significantly faster and scalable query interface
- Goal to be able to pre-process and make available 6-12 months of data from Routeviews
- Release API to research community to allow direct queries to the pre-processed data in addition to the web-based query interface

MRTP

Key Ideas:

- Aggregate BGP UPDATE information from MRT data and generate RPSL-like output summary
- By using RPSL-like format, output can readily be loaded into a RPSL based whois server such as IRRd
- Record reachability times, collector peers, and upstream AS'es in “route:” objects
- By using IRRd, several useful queries can be made – such as searches for more specifics, less specifics, and inverse queries based on origin AS
- Create monthly archives to allow analysis of historical data

Current Status:

- MRTP largely complete, needs some clean up before release
- Generated summaries for Routeviews data back to 2001
- Working on ability to synchronize data in near real-time
- Will be improving IRRd indexing memory utilization so that all db's can be loaded concurrently
  - Currently uses about 2GB of memory for 4 years of data

MRTP "object" examples

```
route:    0.0.0.0/7
origin:   AS13041
beginrch: 2004-12-13 00:57:53Z
endrch:   2004-12-13 01:39:58Z
beginrch: 2004-12-13 01:40:55Z
endrch:   2004-12-13 01:51:23Z
lastann:  2004-12-13 01:40:55Z
rcpeers:  33 (1)
uppeers:  AS4589
source:   RV00-200412


route:    35.0.0.0/8
origin:   AS237
beginrch: 2004-12-01 00:21:59Z
lastann:  2004-12-31 11:27:20Z
rcpeers:  1-39 41 (40)
uppeers:  AS174 AS209 AS3561 AS12956 AS6453 AS2914 AS11537 AS6539
AS3303 AS22335
source:   RV00-200412


peering-set: PRNG-RV00-200412-33
peering:     AS6895 193.149.1.1
updcount:    1525690
source:      RV00-200412
```

AS9121 – Brief facts

AS9121 Turk Telekom – Turkish national telco
Nominally originates about 200 prefixes
Routeviews data shows 60+ AS'es transiting about 500 prefixes via AS9121
Has registered routing policy in RIPE DB
- AS-TTNET as-set in RIPE DB contains 119 AS'es
- aut-num policy is also registered
- import policy for customer peers is “accept ANY” - i.e., no filtering

Major transit peers include
- AS6762 Telecom Italia Sparkle SEA-BONE
- AS1299 TeliaSonera
- AS1239 Sprint
- AS1273 C&W

AS9121 incident on Dec 24 2004

At 09:19 UTC on Dec 24, 2004, AS9121 began re-originating a large number of globally routed prefixes
Peaked at 105,409 prefixes at 9:31 UTC
Lasted until 10:38 UTC – 1 hour, 19 minutes duration
Smaller secondary events also observed
- 11:03 UTC - peak 1238 prefixes - duration 10 minutes
- 19:47 UTC - peak 4579 prefixes - duration 35 minutes

Redistributed primarily via AS6762 (Telecom Italia)
- 106439 unique prefixes seen via AS6762
- Appears they had no filters or prefix limits

Other upstreams had smaller roles
- AS1239 (Sprint) - 5174 prefixes - mostly during final event
- AS1299 (Telia) - 1796 prefixes - max prefix limit of 1000
- AS1273 (C&W) - 162 prefixes - filters?

Total unique prefixes from all peers - 106722

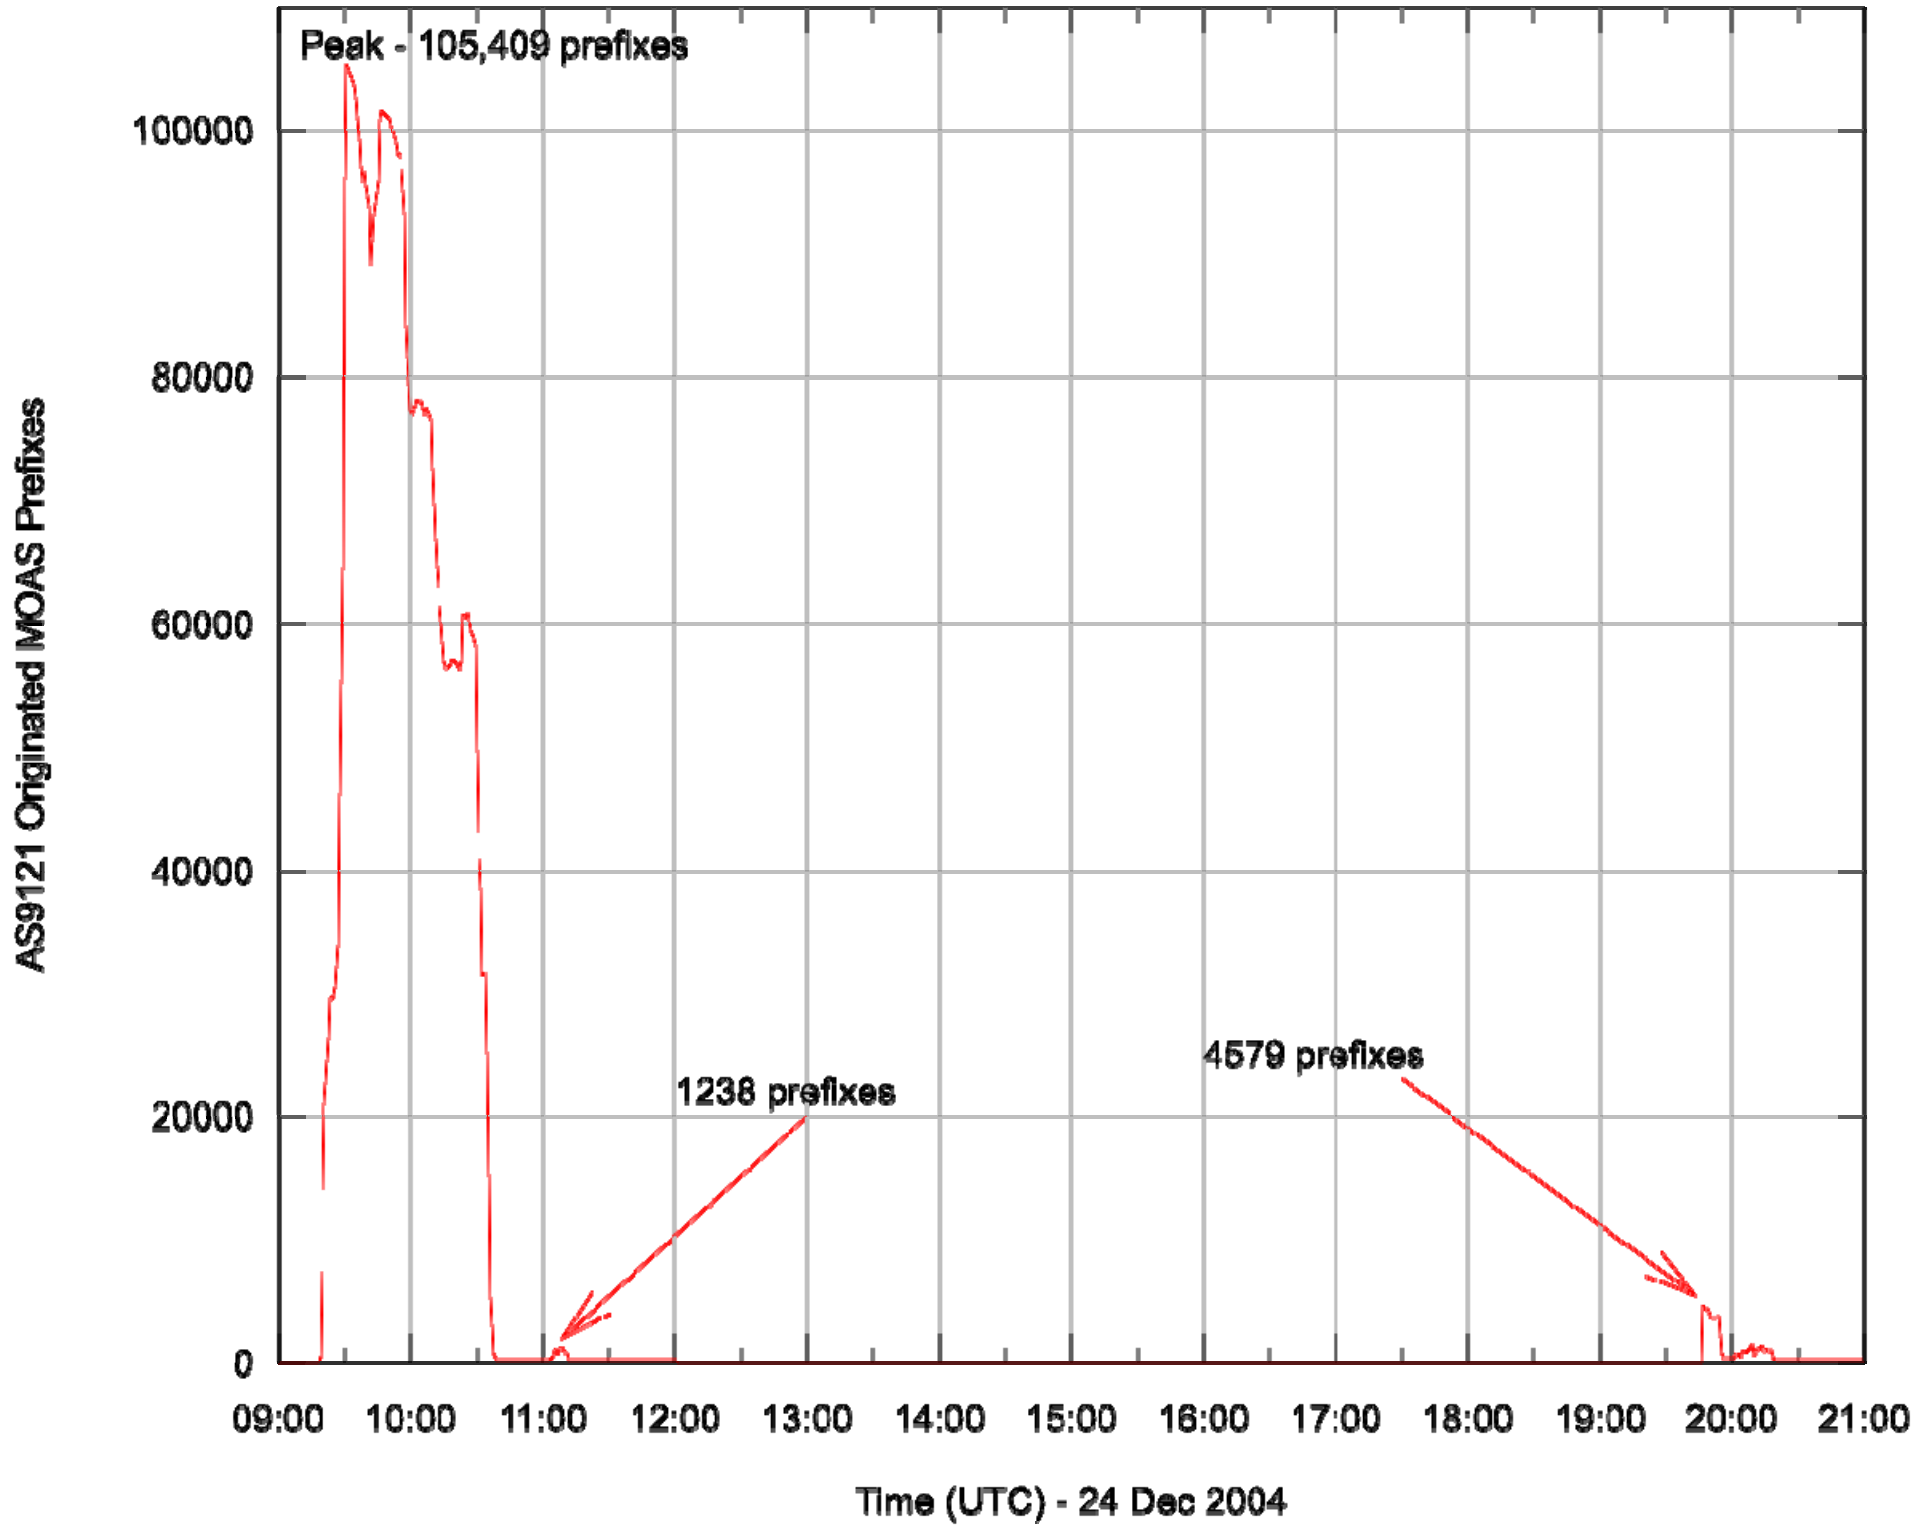
Peak - 105,409 prefixes
1238 prefixes
4579 prefixes
AS9121 Originated MOAS Prefixes
100000
80000
60000
40000
20000
0
09:00
10:00
11:00
12:00
13:00
14:00
15:00
16:00
17:00
18:00
19:00
20:00
21:00
Time (UTC) - 24 Dec 2004

Secondary events (closer look)

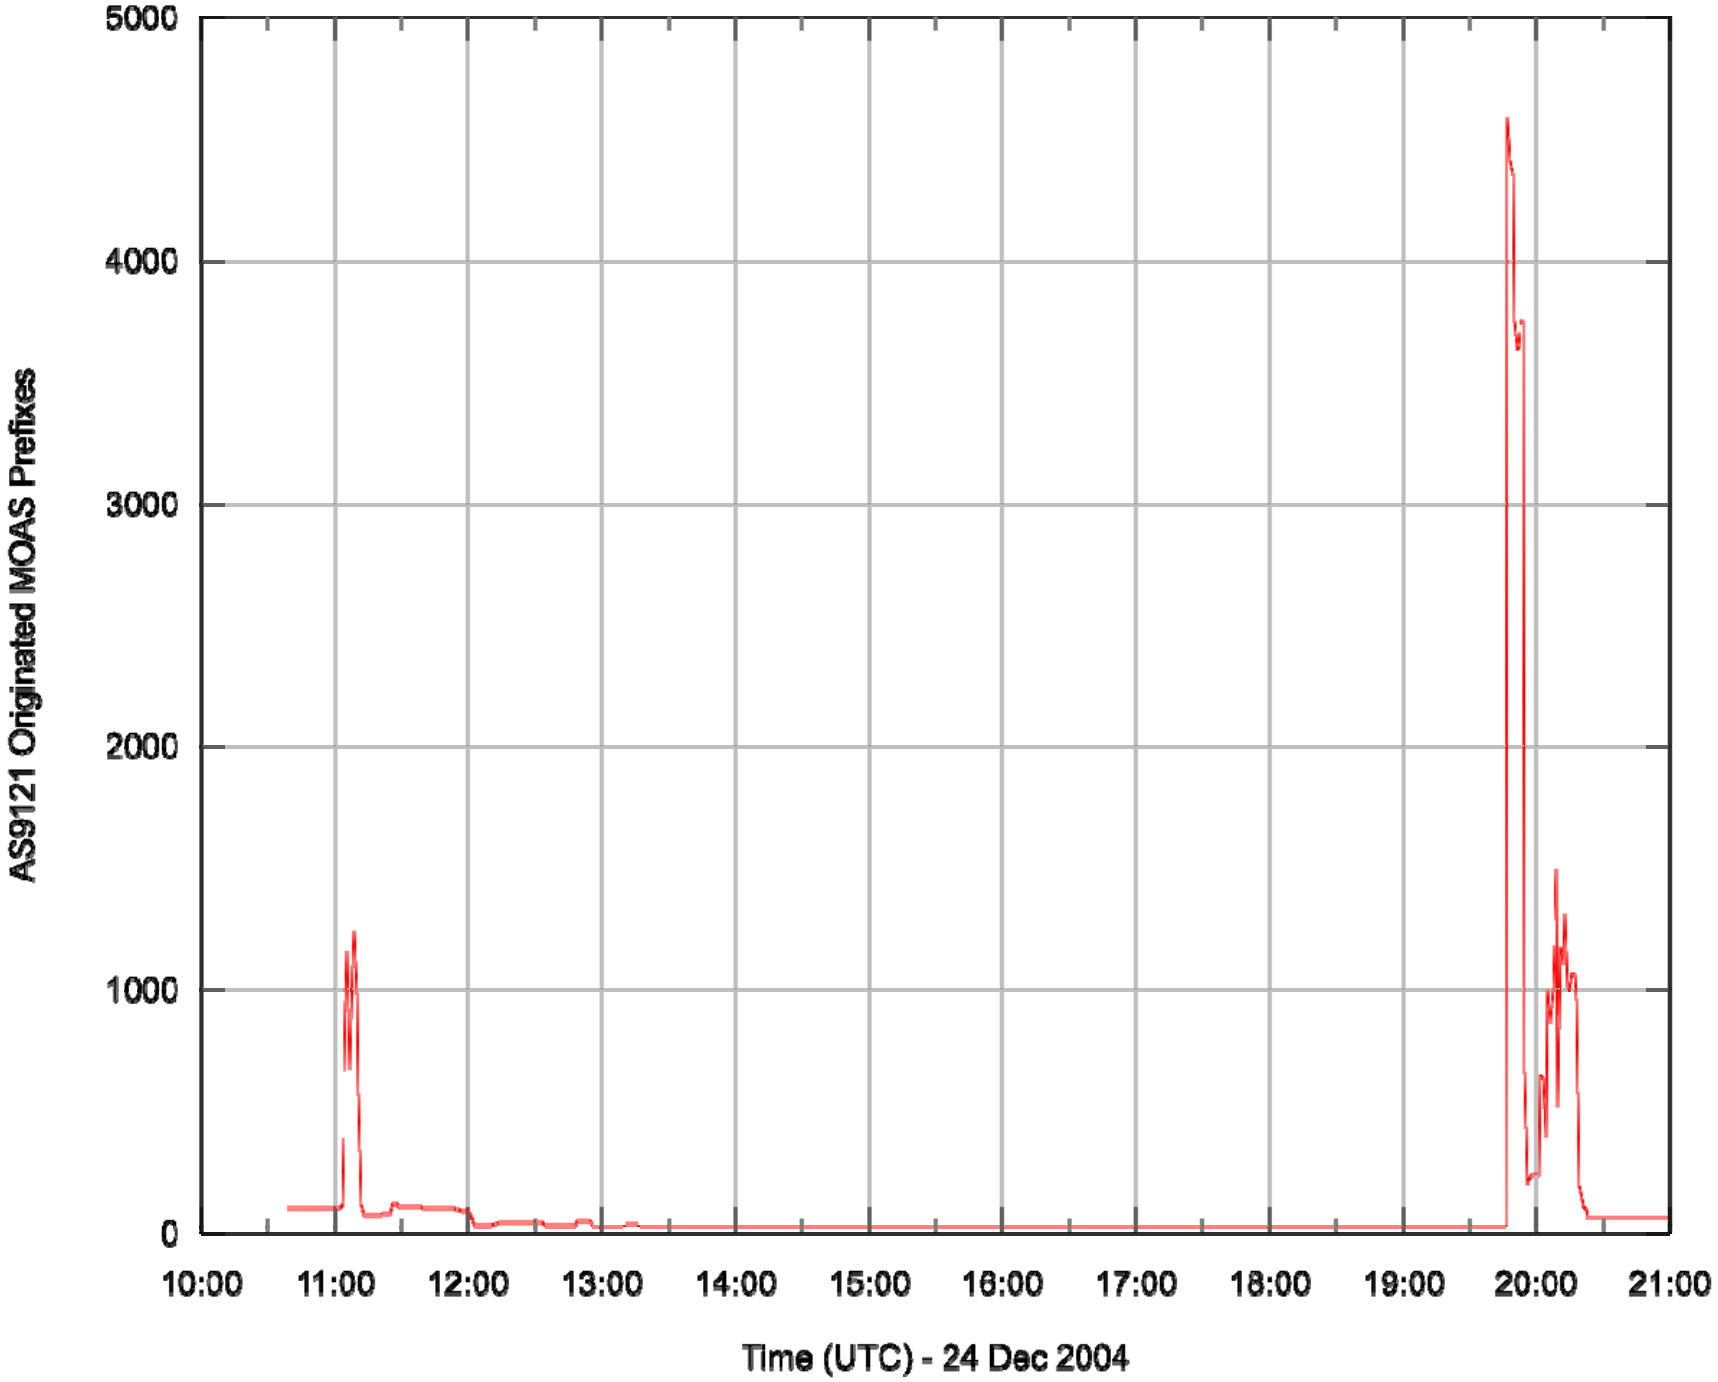

View from Routeviews peers

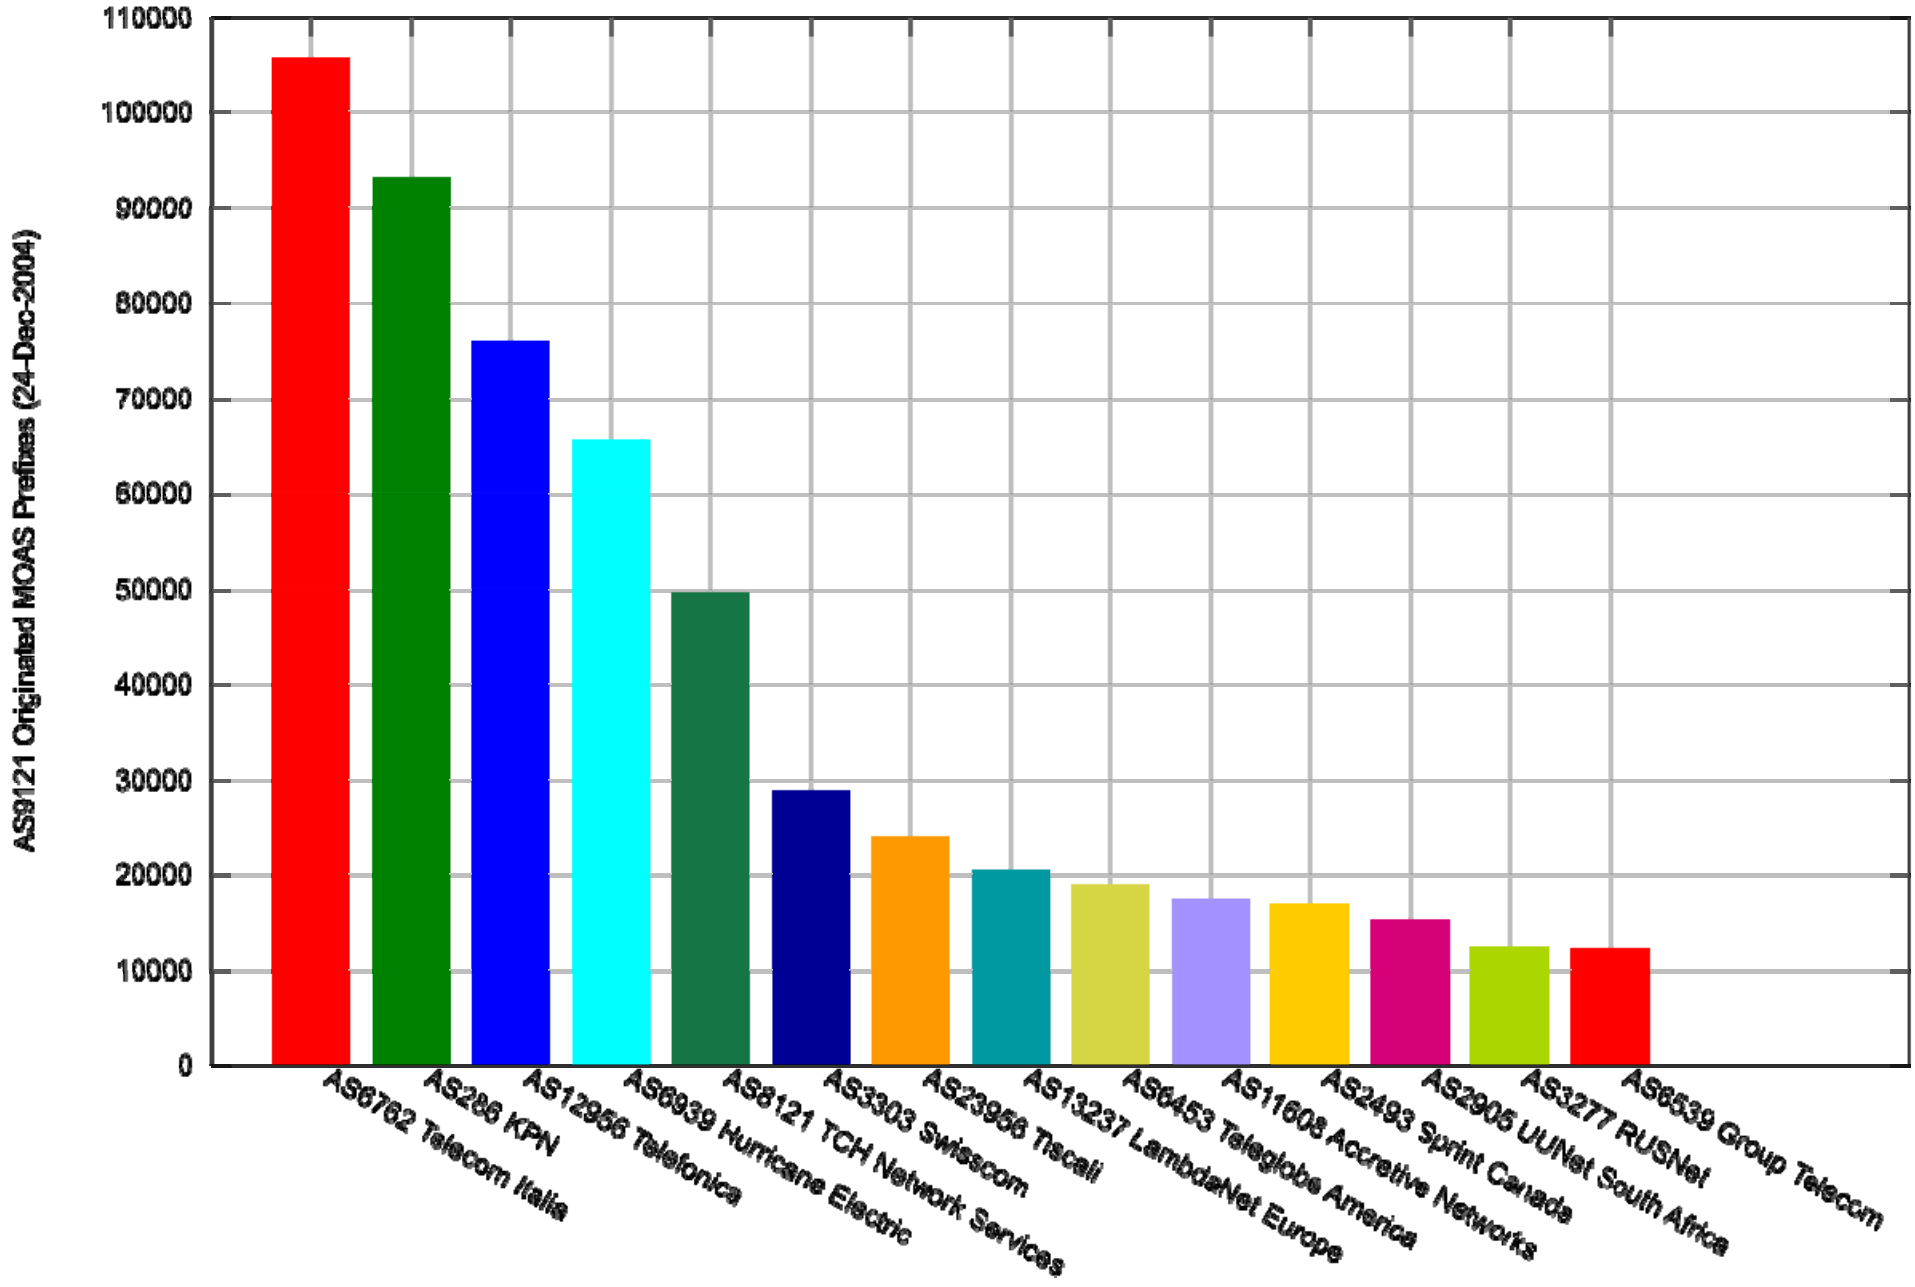

View from Routeviews peers (con'd)