

Using routing protocols for reducing power consumption - a metric-based approach

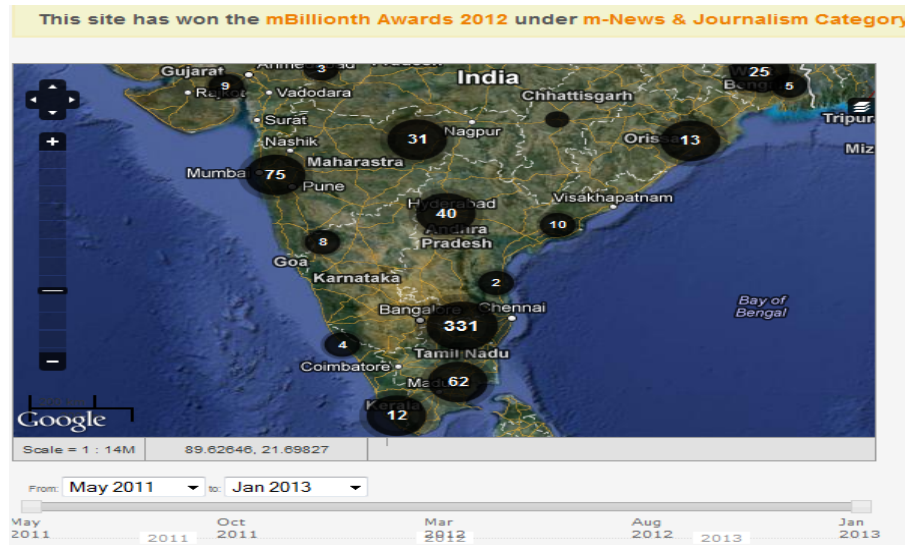
Shankar Raman, Balaji Venkat*, Gaurav Raina
Indian Institute of Technology Madras

Outline

- Motivation
- Power reduction approach
 - Inter-AS,
 - Inter-Area, Intra-Area
- Router power profiles
- Future work
 - Multicast
 - Power reduction trade-off

Power / Internet trends in India

Reported blackouts / brownouts




<http://powercuts.in>

Supply from Grid:

- From 6 to 24 hours
- Thermal power plants

Mobile Internet Penetration

India Internet Traffic by Type, Desktop vs. Mobile, 12/08 – 11/12




<http://ibnlive.in.com/news/mobile-internet-usage-in-india-outstrips-traffic-from-desktops-report/308772-11.html>

Mobile Devices

- High penetration
 - 900 Million Subscribers
 - 600 Million unique
- More Video traffic

Future trends

Power Generation



<http://mospi.nic.in>

2015

Internet

- Video > 50% by 2013
- 9 fold growth
- 53x in 2015 from 2005
- Mobile 28%

Power


- Thermal Power
- Solar (12 hours)
- Wind (Unreliable)

Need:

Protocols that can *automatically re-route to low power consuming devices* during black/brown outs.

Service levels: Audio conference preferable instead of video/audio conference

Motivation: *router power variation*




Picture source : EANTC

Full Chassis MX480 3D

Device level power consumption

- *Base power to full utilization*
 - Approximately 11.6 % change
- *Part of areas and autonomous systems (AS)*
 - Power variation = \sum power variation in router 'i' all routers
- *AS have large variation*
- *Make use of these variations*

Motivation: *variation between routers*

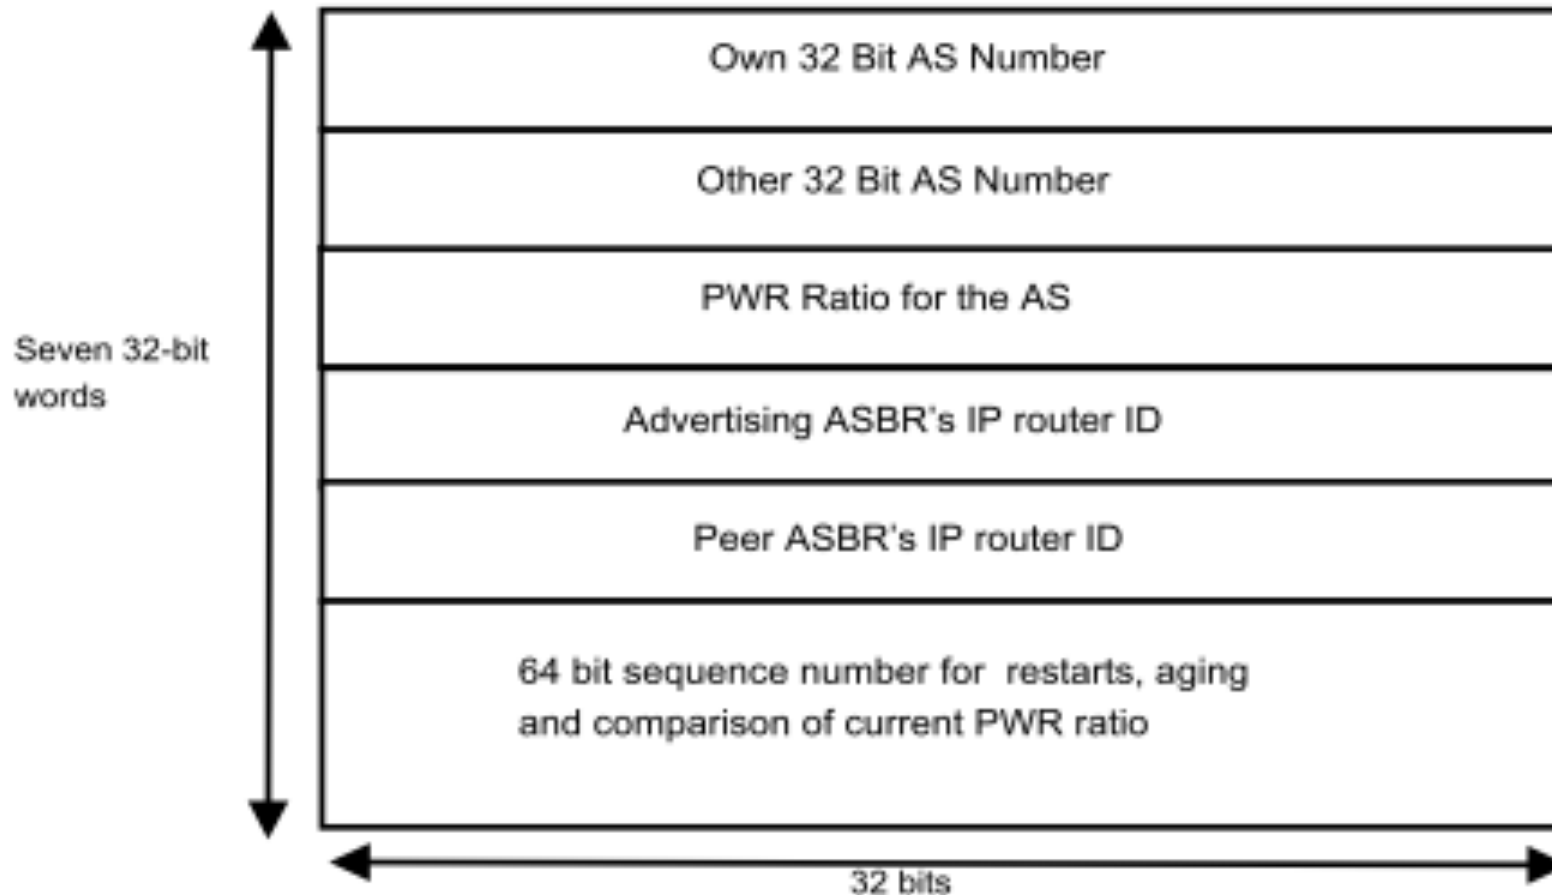


Power profile showing the power consumption in various modes of operation from power up, idle and 70% and 100% link utilization at 1518-byte frames for the switches under test.

Proposed techniques

- *Metric*
 - Consumed power to available bandwidth
 - Consumed power to Multicast Replication Capacity
- *Topology*
 - BGP strands, OSPF/ISIS Link State Database and the OSPF/ISIS TE Database
- *Algorithm*
 - Modified BGP Power Sum Algorithm, Constrained Shortest Path first
- *Traffic Engineering*
 - RSVP TE
- *Hierarchical Approach*
 - Inter-AS and Intra-AS (Inter-Area and Intra-Area)

Example: Using BGP “the how”



Proposed PDU

Modified steps in BGP algorithm

If PWR metric is in use

{

 If there is only one AS path choose this path exit;

 If there are many AS paths

 {

 calculate the PWR based shortest path;

 install it in the routing table ;


 exit;

 }


}

Use the BGP algorithm


Selective low-power path leak (inter-area)



Simulations



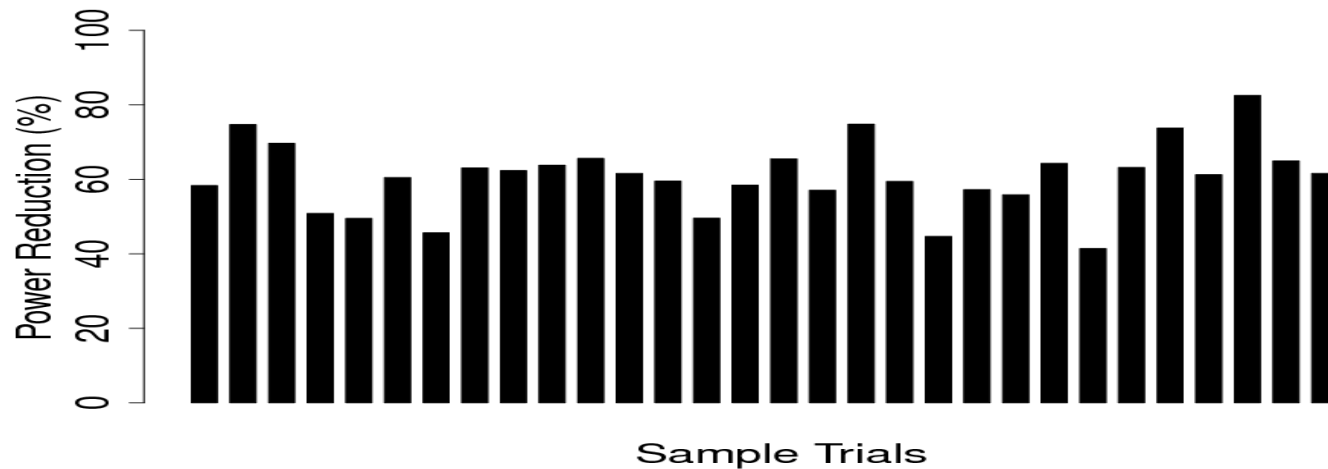
Uniform distribution of PWR



Exponential distribution of PWR

Potential power savings


Not aiming at maximum power savings



Trials = time instants

On average we get power reduction for a given power distribution (in routers, Area, AS)

Router level measurement / profiles



Mix of optical / wired routers from vendors

Case 4: Switch over to Available Power / Available Bandwidth

Future work

- *Create power profile* for areas and AS
 - Cannot be optimal with power points
 - currently used
- Use *metrics* to drive the power reduction
 - Implement for *multicast* as well
- Inter-AS
 - Minimize the *tradeoffs latency/hops, delay-jitter*
 - Research on *buffer sizing and queue management*
 - Increase accuracy of topology discovery

Thank you

Questions ?

Backup

Approach at every level

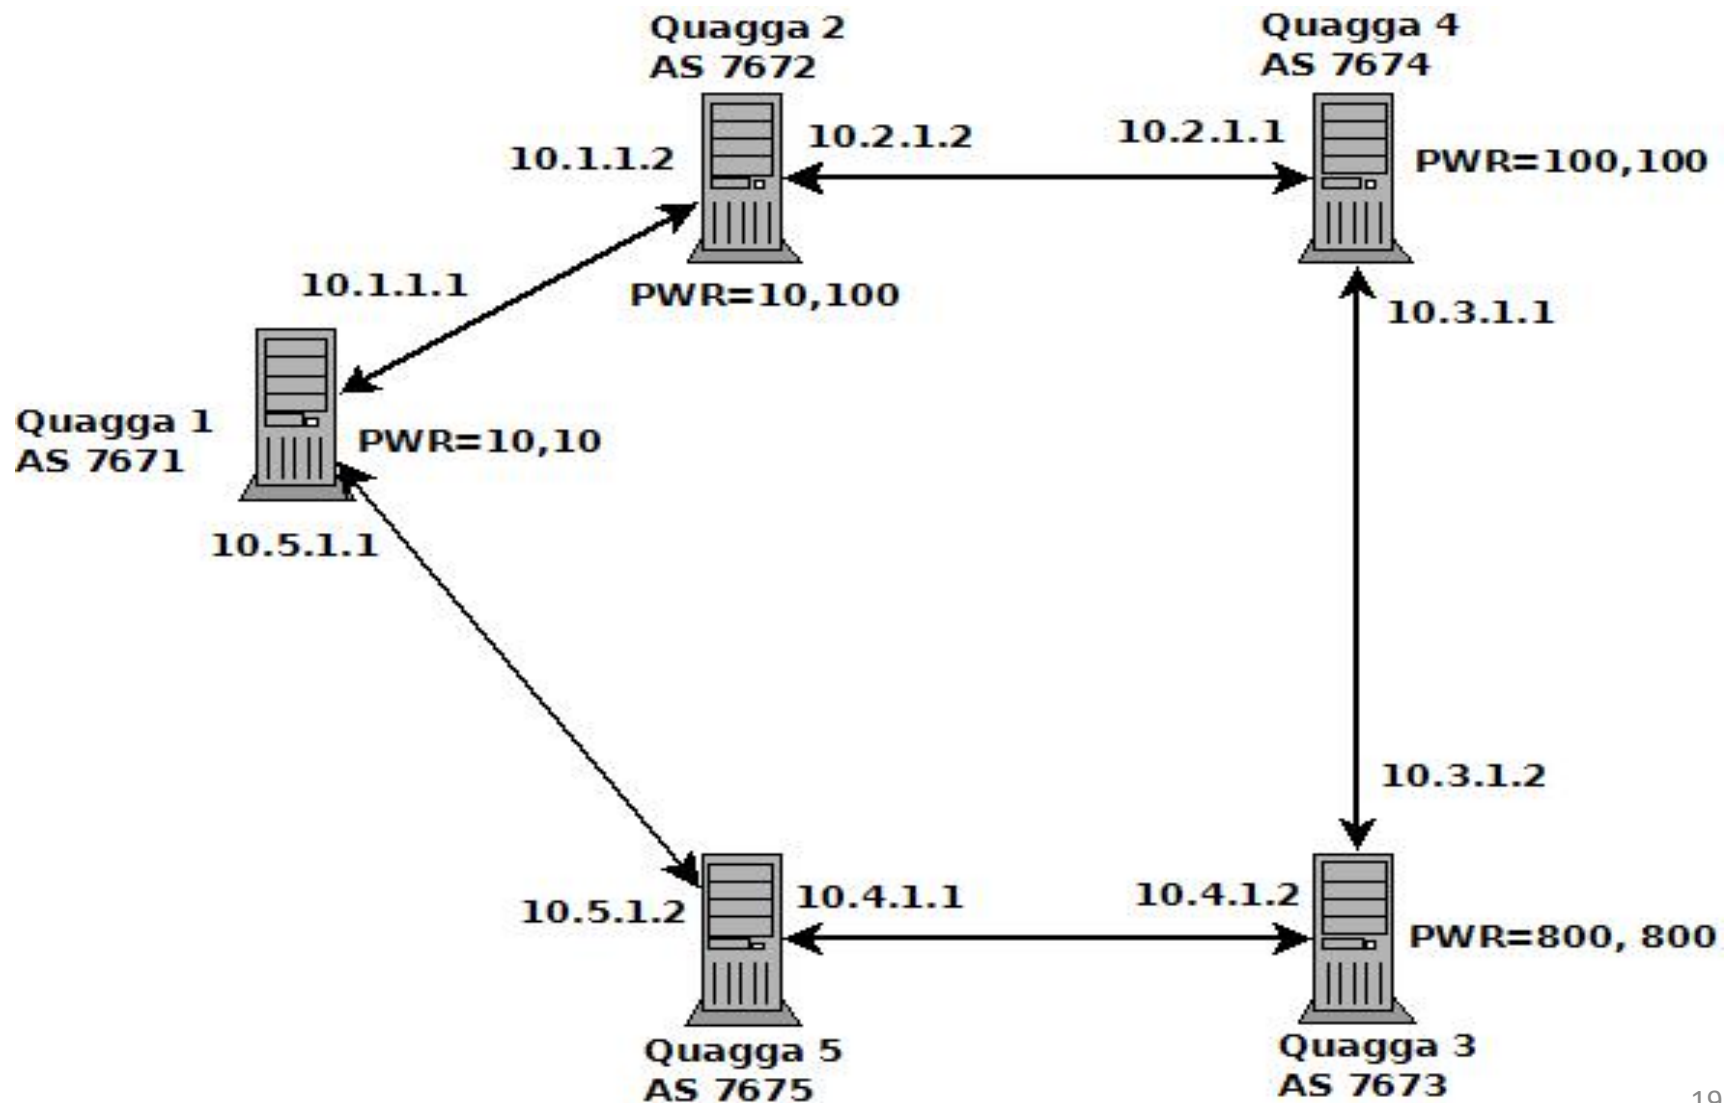
- Metric
 - Consumed Power to Available Bandwidth
 - Consumed Power to Available Replication capacity
- Building the Power based Topology
 - BGP, OSPF, ISIS
 - Use available routing protocols to exchange power information
- Algorithm
 - CSPF, Modified BGP power summation
- Traffic Engineering
 - RSVP TE

Benefits to Operators

- Power reduction techniques can be under the control of the operator or the device vendor.
 - Method can be chosen by the operator.
- Operators get informed about devices that have high power consumption per gigabit of traffic.
 - Little manual intervention needed.
- Protocols take care of the operational power variation
- Can redirect traffic if running on external power sources
- Power variation during traffic load from base power is not differentiated currently.
- Schemes use the hierarchical routing model of the ISPs

Experimental Setup

Quagga Router - Topology

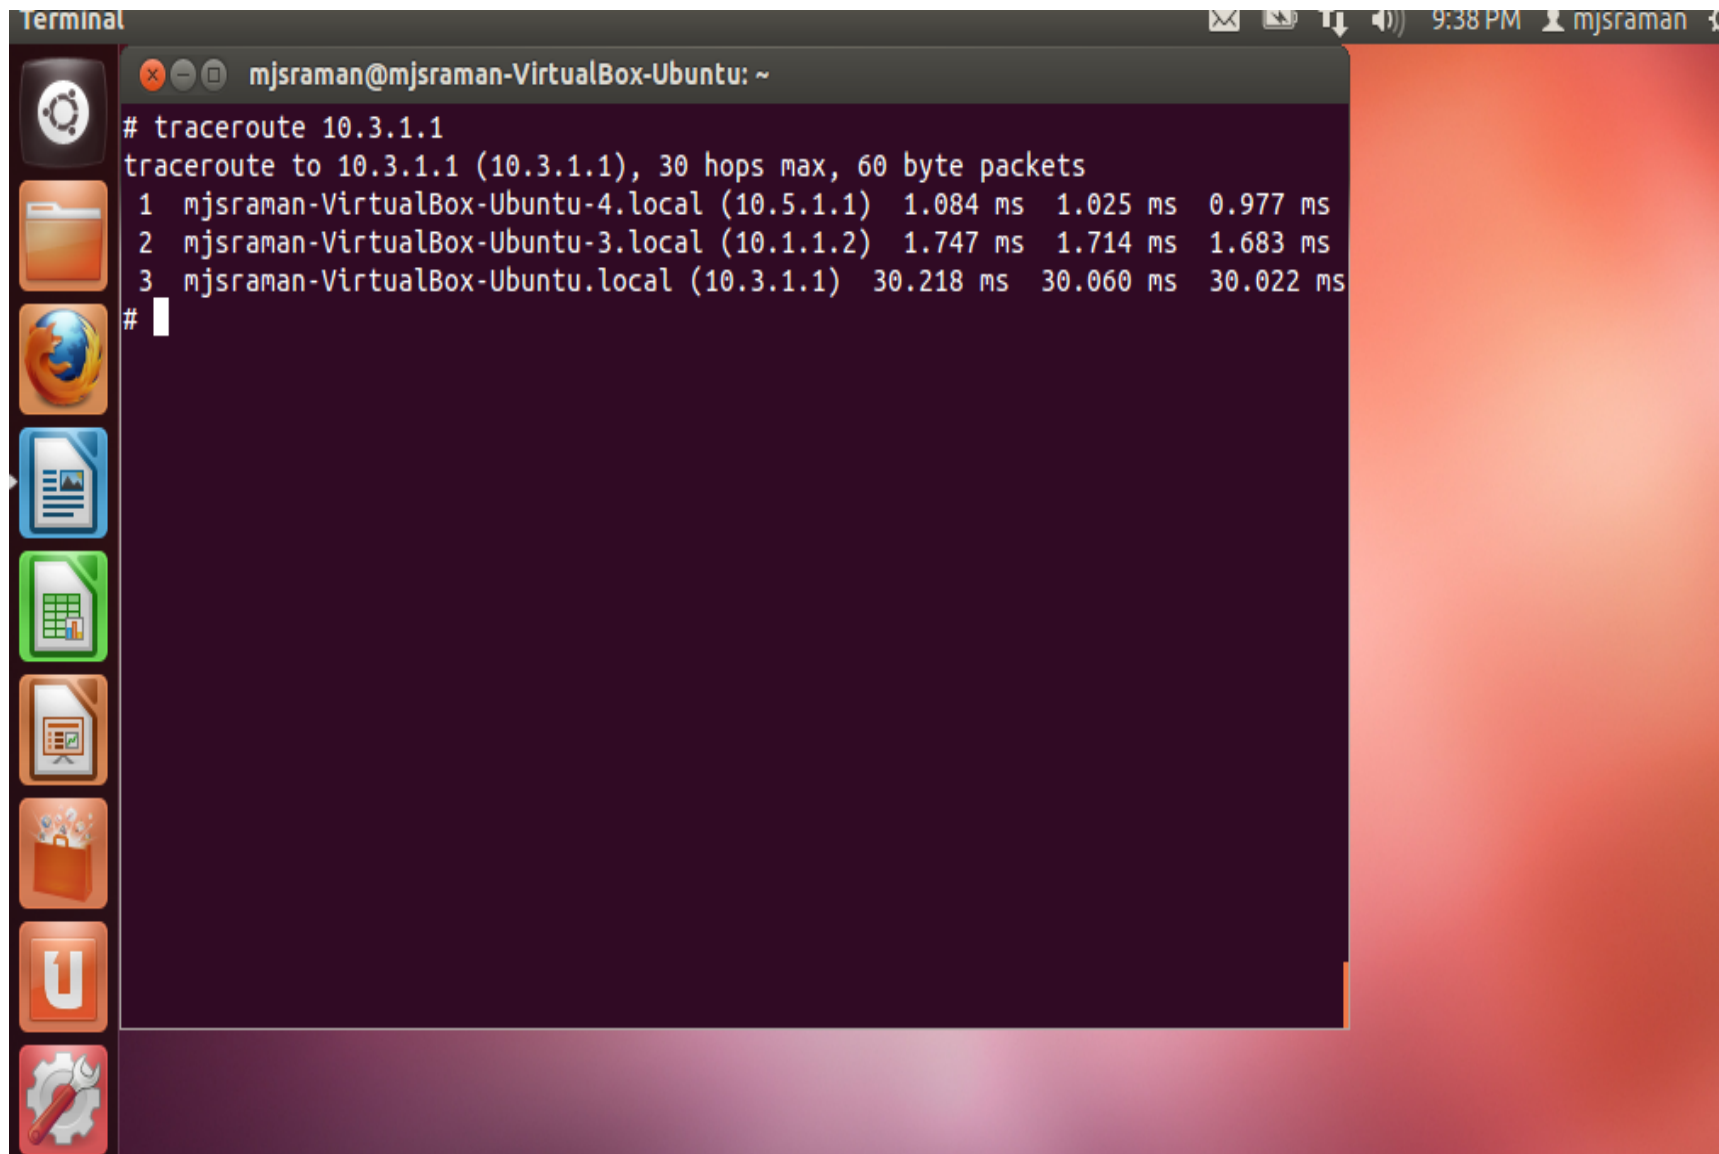


Show ip bgp (routes)

```
Terminal
mjsraman@mjsraman-VirtualBox-Ubuntu: ~
*
10.4.1.2
0
0
0 7673800, i
Total number of prefixes 7
bgpd# sh ip bgp
BGP table version is 0, local router ID is 10.4.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop           Metric LocPrf Weight Path
*> 1.1.1.1/32      10.5.1.1              0           0 7671 7672 7674 767310, 10
, 100, 100, i
*
* 10.1.0.0/16      10.4.1.2              0           0 7673800, i
*> 10.1.0.0/16     10.5.1.1              1           0 767110, ?
*
* 10.2.0.0/16      10.4.1.2              0           0 7673 7674 7672800, 100, 1
0, ?
*> 10.2.0.0/16     10.5.1.1              0           0 7671 767210, 10, ?
*
* 10.3.0.0/16      10.4.1.2              0           0 7673 7674800, 100, ?
*> 10.3.0.0/16     10.5.1.1              0           0 7671 7672 767410, 10, 100
, ?
*
* 10.4.0.0/16      10.4.1.2              1           0 7673800, ?
* 10.4.0.0/16      10.4.1.2              1           0 7673800, ?
*>
* 10.5.0.0/16      0.0.0.0              1          32768 0 (null)?
* 10.5.0.0/16      10.5.1.1              1           0 767110, ?
*>
* 10.15.1.1/32     0.0.0.0              1          32768 0 (null)?
*> 10.15.1.1/32    10.5.1.1              0           0 7671 7672 7674 767310, 10
, 100, 100, i
*
10.4.1.2
0
0
0 7673800, i
Total number of prefixes 7
bgpd#
```

Traceroute from AS 7675



A terminal window titled "terminal" with a dark background and a light-colored text. The window shows the execution of a traceroute command. The output displays three hops, each with the destination IP, the hop number, and three round-trip time measurements in milliseconds. The first two hops are local to the virtual machine, and the third hop is the destination IP 10.3.1.1.

```
mjsraman@mjsraman-VirtualBox-Ubuntu: ~  
# traceroute 10.3.1.1  
traceroute to 10.3.1.1 (10.3.1.1), 30 hops max, 60 byte packets  
1 mjsraman-VirtualBox-Ubuntu-4.local (10.5.1.1) 1.084 ms 1.025 ms 0.977 ms  
2 mjsraman-VirtualBox-Ubuntu-3.local (10.1.1.2) 1.747 ms 1.714 ms 1.683 ms  
3 mjsraman-VirtualBox-Ubuntu.local (10.3.1.1) 30.218 ms 30.060 ms 30.022 ms  
#
```

Hierarchical Approach with Protocols

- AS level topology – intra-AS power reduction
 - BGP
 - draft-mjsraman-panet-bgp-power-path-01
 - Without BGP
 - draft-mjsraman-panet-inter-as-psp-01
- Intra-AS
 - Inter-area, Intra-area
 - draft-mjsraman-panet-ospf-power-topo-00
 - draft-mjsraman-panet-intra-as-psp-te-leak-00
- Multicast, ECMP
 - draft-mjsraman-panet-pce-power-mcast-replic-00
 - draft-mjsraman-panet-pim-power-00
 - draft-mjsraman-panet-ecmp-redirect-power-repl-cap-00