

Collection and Characterization of BCNET BGP Traffic

Sukhchandan Lally, Tanjila Farah, Rajvir Gill, Ravinder Paul,
Nabil Al-Rousan, and Ljiljana Trajković
Simon Fraser University
Vancouver, British Columbia, Canada

E-mail: {lally, tfarah, rajvirg, rpa28, nalrousa, ljilja}@sfu.ca
<http://www.ensc.sfu.ca/~ljilja/cnl>



Roadmap

- Border Gateway Protocol
- BCNET packet capture
- BCNET traffic
- Views of BCNET Traffic using Wireshark
- Preliminary analysis
- Conclusions



What is a Border Gateway Protocol (BGP)?

- BGP is de-facto Inter Autonomous System routing protocol
- Peer routers exchange four types of messages: open, update, notification, and keepalive
- BGP utilizes a path vector algorithm called the best path selection algorithm to select the best path
- BGP routing tables are publicly available and may be retrieved from the Route Views and Réseaux IP Européens (RIPE)



Roadmap


- Border Gateway Protocol
- BCNET packet capture
- BCNET traffic
- Views of BCNET Traffic using Wireshark
- Preliminary analysis
- Conclusions



BCNET Packet Capture

- BCNET is the hub of advanced telecommunication network in British Columbia, Canada that offers services to research and higher education institutions
- BCNET transits have two service providers with 10 Gbps network links and one service provider with 1 Gbps network link
- Optical Test Access Point (TAP) splits the signal into two distinct paths and the signal splitting ratio from TAP may be modified
- The Data Capture Device (NinjaBox 5000) collects the real-time data (packets) from the traffic filtering device


BCNET Packet Capture



Physical overview of the BCNET packet capture

Net Optics Director 7400


- It is used for BCNET traffic filtering
- It directs traffic to monitoring tools such as NinjaBox 5000 and FlowMon



Net Optics Director application diagram

Endace Data Acquisition and Generation Card

- Endace Data Acquisition and Generation (DAG) 5.2X card resides inside the NinjaBox 5000
- It captures and transmits traffic and has time-stamping capability
- DAG 5.2X is a single port Peripheral Component Interconnect Extended (PCIe) card and is capable of capturing on average Ethernet traffic of 6.9 Gbps



Network monitoring and analyzing Endace card




Roadmap

- Border Gateway Protocol
- BCNET packet capture
- **BCNET traffic**
- Views of BCNET Traffic using Wireshark
- Preliminary analysis
- Conclusions

BCNET Traffic

- The BCNET network is high-speed fiber optic research network
- British Columbia's network extends to 1,400 kilometres and connects Kamloops, Kelowna, Prince George, Vancouver, and Victoria



Real time network usage by BCNET members.



Roadmap

- Border Gateway Protocol
- BCNET packet capture
- BCNET traffic
- Views of BCNET Traffic using Wireshark
- Preliminary analysis
- Conclusions

Views of BCNET Traffic using Wireshark

- Wireshark is an open source packet analyzer that captures network packet data from a network interface and displays those packets with detailed protocol information
- Wireshark provides comprehensive statistics such as the summary of traffic collected, input/output graphs, protocol hierarchy, and endpoints

```
⊞ Frame 298702: 160 bytes on wire (1280 bits), 160 bytes captured (1280 bits)
⊞ Ethernet II, Src: JuniperN_d3:80:d4 (00:23:9c:d3:80:d4), Dst: JuniperN_8e:d0:00 (00:1f:12:8e:d0:00)
⊞ Internet Protocol, Src: 72.51.24.189 (72.51.24.189), Dst: 72.51.24.190 (72.51.24.190)
⊞ Transmission Control Protocol, Src Port: bgp (179), Dst Port: 58268 (58268), Seq: 22708555, Ack: 67964, Len: 74
⊞ Border Gateway Protocol
  ⊞ UPDATE Message
    Marker: 16 bytes
    Length: 74 bytes
    Type: UPDATE Message (2)
    Unfeasible routes length: 0 bytes
    Total path attribute length: 43 bytes
    ⊞ Path attributes
      ⊞ ORIGIN: IGP (4 bytes)
      ⊞ AS_PATH: 13768 20161 19053 (17 bytes)
      ⊞ NEXT_HOP: 72.51.24.189 (7 bytes)
      ⊞ COMMUNITIES: 13768:64995 13768:65002 13768:65507 (15 bytes)
    ⊞ Network layer reachability information: 8 bytes
      ⊞ 199.27.216.0/21
      ⊞ 199.27.223.0/24
```

General Wireshark view of the collected traffic



BCNET Traffic Summary

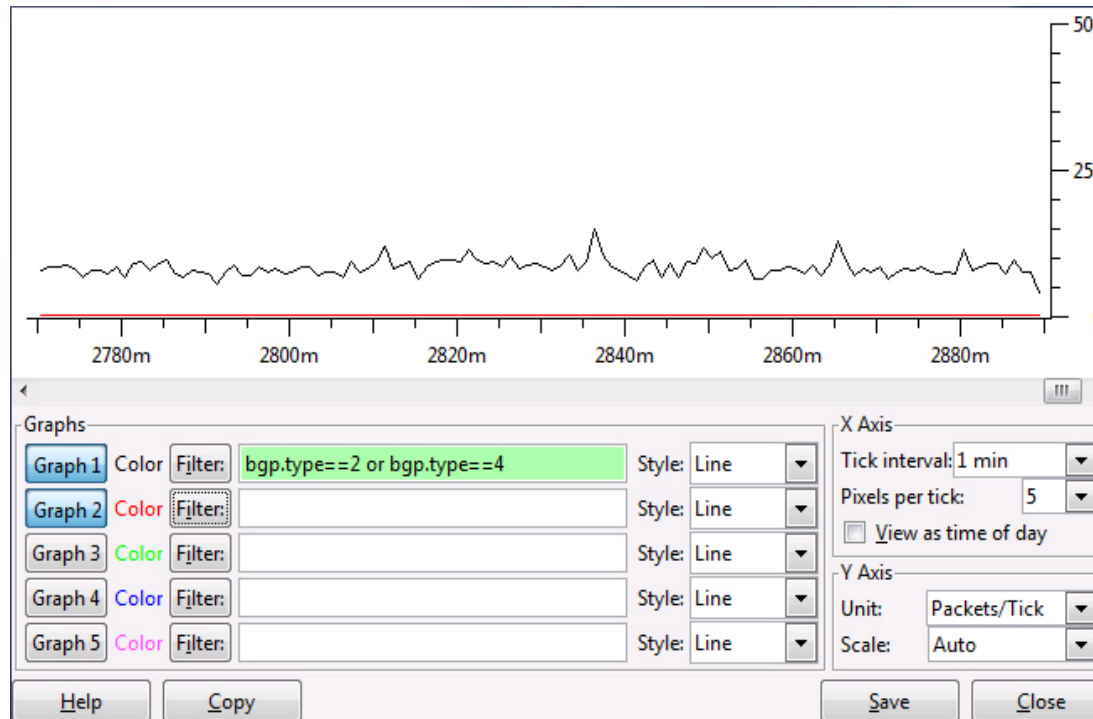
- There were 511,820 packets collected over the period of 48 hours

Time			
First packet:	2010-12-20 14:56:31		
Last packet:	2010-12-22 15:06:05		
Elapsed:	02 days 00:09:34		
Display			
Display filter:	bgp.type==2 or bgp.type==4		
Ignored packets:	0		
Traffic	Captured	Displayed	Marked
Packets	511820	260639	0
Between first and last packet	173374.253 sec	173374.154 sec	
Avg. packets/sec	2.952	1.503	
Avg. packet size	192.046 bytes	305.514 bytes	
Bytes	98292937	79628747	
Avg. bytes/sec	566.941	459.288	
Avg. MBit/sec	0.005	0.004	

Summary of BCNET traffic collected

BCNET Traffic Input-Output Graphs

- BCNET Traffic Input-Output Graphs define up to five filters
- The number of samples is limited to 100,000



Input-Output graph of the packets captured



BCNET Traffic Protocol Hierarchy statistics

- Each protocol has its statistical value (row) consisting of protocol's name, the percentage of protocol packets relative to total number of packets captured, number of packets, and number of bytes
- From 511,820 packets, 260,639 (50.9%) are BGP packets, 257,285 (50.3%) are TCP ACK packets, and 6,104 (1.2%) are piggyback ACKs

Protocol hierarchy of the captured packets

Protocols	Packets %	Packets	Bytes
Ethernet/IP/TCP	100	511,820	98,292,937
BGP	50.92	260,639	79,628,747



BCNET Network Endpoints

- They are the source and destination addresses of a specified protocol layer
- Endpoints of the six BCNET transit exchanges (BGP peers) are captured
- There are various TCP connection statistics for each IP address of a BGP peer

Network Endpoints

Address	Port	Packets	Bytes	Tx Bytes	Rx Bytes
72.51.24.189	bgp	401721	70836354	55894998	14941356
72.51.24.190	58268	401721	70836354	14941356	55894998
64.251.87.209	bgp	70069	14996289	12426684	2569605
64.251.87.210	62844	70069	14996289	2569605	12426684
206.108.83.66	bgp	40030	12460294	1500045	10960249
206.108.83.70	51899	40030	12460294	10960249	1500045

BCNET Traffic Service Response Time

- It is defined as the time between a request and the corresponding response
- The flowgraph of the BGP peers includes the source address, destination address, TCP port number, TCP message (ACK), and type of the BGP message (open, update, notification, keepalive)



August 25, 2011




Roadmap

- Border Gateway Protocol
- BCNET packet capture
- BCNET traffic
- Views of BCNET Traffic using Wireshark
- **Preliminary analysis**
- Conclusions

Preliminary Analysis of BCNET Traffic


- Various factors such as link speed, propagation delay, window size, link reliability, and congestion of network and intermediate device affect the throughput of TCP
- The throughput graph illustrates that the average throughput of the collected data is 177.1 packets/min and that the maximum throughput is 645 packets/min




TCP throughput

TCP window size and TCP RTT

- TCP window size of the BCNET traffic for 200 samples
- TCP RTT
 - TCP RTT of the BCNET traffic with an average of 11.7 ms
 - The RTT standard deviation is 7.19 ms, 2.75 ms for the sample and estimated RTT



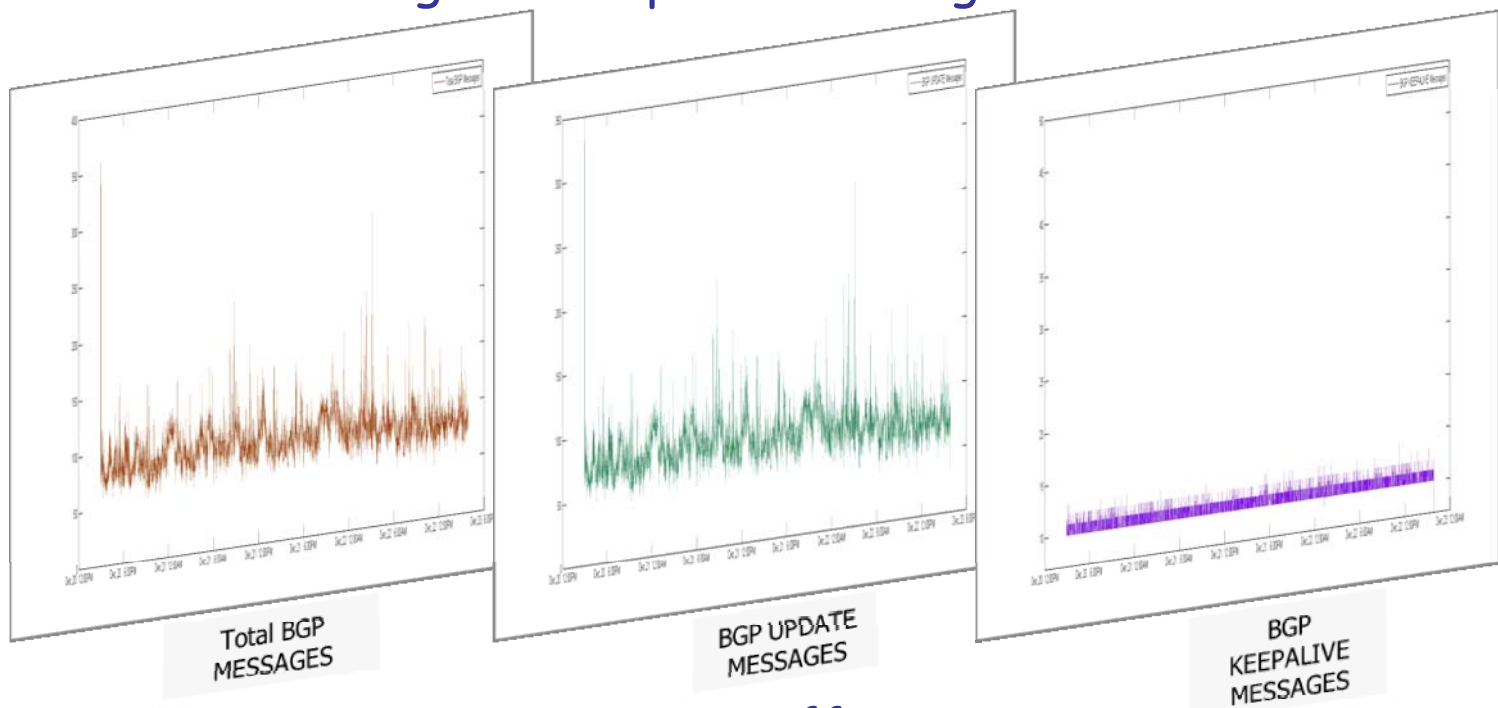
TCP window size



TCP RTT

BGP Messages

- Sample of display filters are: `bgp.type`, `bgp.next_hop`, `bgp.origin`, `bgp.local_pref`, `bgp.community_as`, `bgp.as_path`, and `bgp.multi_exit_disc`
- In the collected BGP traffic, 88% are BGP update messages and the remaining are keepalive messages




August 25, 2011

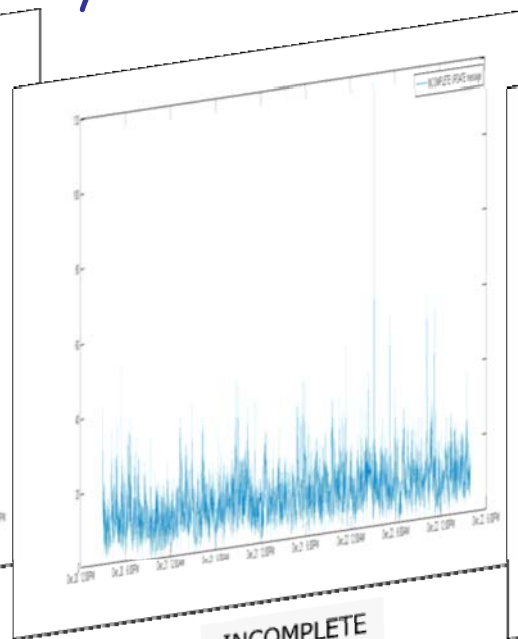
BGP traffic

BGP Message attributes

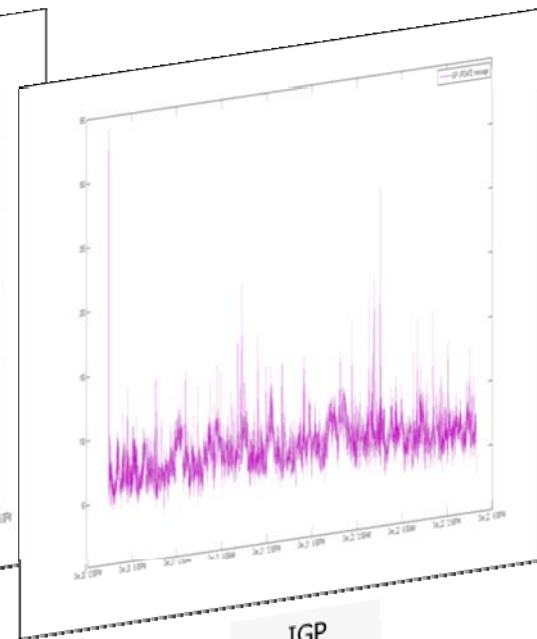
- BGP origin which defines the origin of the path may have three values: Interior Gateway Protocol (IGP), Exterior Gateway Protocol (EGP), or Incomplete
- The EGP, Incomplete, and IGP messages account for 0.003%, 13.84%, and 85.82% of the total number of BGP update messages, respectively



EGP UPDATE
MESSAGES



INCOMPLETE
UPDATE
MESSAGES



IGP
UPDATE
MESSAGES

August 25, 2011



Roadmap

- Border Gateway Protocol
- BCNET packet capture
- BCNET traffic
- Views of BCNET Traffic using Wireshark
- Preliminary analysis
- **Conclusions**



Conclusions

- The collected data will be used to analyze performance of the BGP protocol and the effect of route flaps and parameters such as the minimal route advertisement interval (MRAI)
- BGP traffic data collected from BCNET will be compared to Internet topologies generated from the publicly available Route Views and RIPE datasets



References

- Y. Rekhter, T. Li, and S. Hares, "A border gateway protocol 4 (BGP-4)," *IETF RFC 1771*.
- Autonomous System Numbers [Online]. Available: <http://www.iana.org/assignments/as-numbers>.
- BGP Best Path Selection Algorithm [Online]. Available: <http://www.cisco.com/en/US/tech/tk365>.
- L. Gao, "On inferring autonomous system relationships in the Internet," *IEEE/ACM Trans. Networking*, vol. 9, no. 6, pp. 733–745, Dec. 2001.
- BGP capture datasets [Online]. Available: <http://archive.routeviews.org>.
- Réseaux IP Européens [Online]. Available: <http://www.ripe.net/ris>.
- BGP Routing Table Analysis Reports [Online]. Available: <http://www.potaroo.net/bgp/>.
- BCNET [Online]. Available: <http://www.bc.net>.
- Data Monitoring Switch [Online]. Available: <http://www.netoptics.com/products/director>.



References

- Wireshark [Online]. Available: <http://www.wireshark.org>.
- Welcome to DAG [Online]. Available: <http://www.endace.com>.
- BCNET Traffic Map [Online]. Available: <https://www.bc.net/atlconf/display/Network/BCNET+Traffic+Map>.
- BGP Case Studies [Online]. Available: <http://www.cisco.com/application/pdf/paws/26634/bgp-toc.pdf>.
- W. Shen and Lj. Trajković, "BGP route flap damping algorithms," in Proc. SPECTS 2005, Philadelphia, PA, July 2005, pp. 488–495.
- N. Lasković and Lj. Trajković, "BGP with an adaptive minimal route advertisement interval," in Proc. 25th IEEE Int. Performance, Computing, and Communications Conference, Phoenix, AZ, Apr. 2006, pp. 135–142.
- Lj. Trajković, "Analysis of Internet topologies," IEEE Circuits and Systems Magazine, vol. 10, no. 3, pp. 48–54, Third Quarter 2010.