

Mapping the Internet

Horms (Simon Horman)
horms@valinux.com
Senior Software Engineer
VA Linux Systems

February 2001
http://vergenet.net/~horms/inet_map/

Live in Northern California once
but leave before it makes you soft.
Live in New York once
but leave before it makes you hard.

Mary Schmich, Baz Luhrmann
Everybody's Free (To Wear Sunscreen)

Licence

This paper is copyright 2000 Simon Horman and is released under the terms of the GNU General Public Licence, a copy of which is distributed with the source for this document.

All trademarks and software are the property of their respective owners.

Abstract

The internet is a large, complex interconnection of different networks. Other, simpler networks such as suburban rail systems are mapped to give users omniscient knowledge. Wouldn't it be nice to do the same to the Internet?

1 Introduction

BGP[4][1][3] is an Exterior Routing protocol designed to allow networks to communicate information on how to reach hosts within themselves and connected networks. This is the protocol that is used on the internet to allow traffic to traverse to globe or more accurately the networks that form the internet. BGP is adaptive to changes in network topology that may be schedule or occur due to equipment or link failures.

BGP makes information about the networks that traffic to a given *prefix*¹ will travel through. This information is typically used for debugging network problems. Creative people such as myself like to this information for other purposes such as load balancing traffic across the internet[2].

Given that BGP provides information on what networks traffic will travel through it would seem that this could be used to construct a map of the internet. After all, a map of the internet could be represented as the links between the networks that form the internet.

This was the topic of conversation I had on a recent long-haul flight. I thought that I would investigate it further.

2 BGP and Inter-Network Links

A host running a BGP session to one or more peers stores information about how to get to prefixes which is is connected to, directly or indirectly. This includes information on the networks that traffic will travel through to get to the prefix².

Suppose there are three networks A, B and C interconnected by links AB, AC and BC. Each network runs BGP to its neighbouring networks over each link. The BGP session running at A has information from B about how to get to B. That is A knows that to get to an address in B it can send traffic across AB. A also has information from B about how to get to C, again by sending traffic across AB, who will then forward it directly to C. A can deduce from this information that there is a link BC.

¹A prefix is a network address and a bit mask. The mask is used to differentiate the network address from hosts within the network. In the context of IPv4 on the internet network addresses are IP addresses and bit masks may be given as the number of significant bits in the bitmask. For example 10.0.0.0/30 has a network address of 10.0.0.0 and has 30 significant bits in the bitmask. That is the network covers the addresses 10.0.0.0–10.0.0.3.

²It is of note that if asymmetric routing is being used then the return path for packets may well travel through different networks.




Figure 1: Inter-Network Diagram — Three Interconnected Networks

It follows that each of A, B and C have direct or indirect knowledge of all the inter-network links AB, AC and BC. It would seem that BGP can indeed be used to create a map of inter-connected networks and indeed a map of the internet.

3 Hidden Links




Figure 2: Inter-Network Diagram — A Leaf Network

Suppose an additional network D is added which is connected only to C via link CD; D is a leaf network in the inter-network. Given that D has a BGP session established with C it will learn how to get to each of A, B and C. C will send D information on what it considers to be the least cost way to get to each of these networks. Assuming that all links and networks are available and ignoring the possibility of links being weighted, this information will be follows:

Destination	Path
A	C A
B	C B
C	C

D has direct knowledge of the link CD and can deduce that links BC and AC are present. But as the link AB is not the least cost way to get to any network it does not appear in any path. As it is not in any path D has no knowledge of this link.

If link AC was to fail then C would inform D that the least cost path to A is C B A. At this time D would be able to deduce that the link AB exists. But the problem still remains that at any given time it is not necessarily possible to see all links present on the inter-network. Thus, omniscient knowledge of the inter-network cannot be obtained and a map of the internet, showing all inter-network links cannot be constructed using this information alone.

4 Conclusion

One can think of the information supplied by BGP as a set of vectors running through the inter-network. To be sure that all vectors and, hence, all links were known then information would be needed from all routers running BGP. Clearly, on the internet this is not practical.

It is possible to construct a map of the links that can be seen from a given BGP session. This in itself is quite useful, it would be the internet as a given network sees it. It may be possible to use this information in conjunction with other sources of information to construct a more thorough map. This will be left for discussion on another plane trip.

References

- [1] Avi Freedman. Bgp routing part i: Bgp and multi-homing, 1997. <http://www.netaxs.com/~freedman/bgp/bgp.html>.
- [2] Simon Horman. Global load balancing (using bgp to take over the world), December 2000. http://supersparrow.org/ss_paper/.
- [3] Christian Huitema. *Routing In The Internet*. Prentice Hall, 2 edition, 2000.
- [4] Y. Rekhter and T Li. Rfc 1773: A border gateway protocol 4 (bgp-4), 1995.