

# Problem Set 1: Interdomain routing

Instructor: V. Arun  
CS 653, Fall 2010  
Oct 5, 2010

## 1. (Short answer questions):

- If the Stable Paths Problem is solvable, then no dispute wheels exist. True or false?
- If SPVP converges to a unique stable state, then no dispute wheels exist. True or false?
- Is the provider<peer<customer preference an export policy or an import policy?
- Let a *size sequence*  $x_1, x_2, \dots, x_n$  denote the sizes of the ASes in an AS path of length  $n$ . Assume that two ASes have a peering relationship iff they are within a factor of 1.5 of each other in size. Which of the following size sequences can not correspond to a valley-free AS path: (i) [12, 17, 43, 102, 4], (ii) [3, 6, 9, 12], (iii) [1, 2, 4, 8, 1]?
- Can route flap damping prevent oscillations such as in the bad gadget? Explain.
- In a complete AS graph of size  $n$ , how long will BGP take to converge (assuming no dispute wheels exist)? Explain.
- In a complete AS graph of size  $n$ , suppose all nodes restricted paths to be of length at most  $k < n$ . What is the worst-case message processing overhead of BGP just after a routing event if no rate-limiting timers are used? Explain.
- In a complete graph of size  $n$ , if nodes were doing *distance vector* routing based on positive link weights, what is an upper bound on the number of rounds for convergence assuming that nodes start with empty routing tables and neighboring nodes exchange distance vectors once every round.
- In the previous question on distance vector routing, what is an upper bound on the number of rounds for re-convergence just after a routing event?

## 2. (Playing with BGP table dumps):

To answer this question, you will need to download the Routeviews routing table dump at:

<http://www.cs.umass.edu/~arun/653/PS/oix-full-snapshot-latest.dat.gz>

The file contains a Cisco BGP routing table snapshot taken at the Oregon Routeviews site <http://www.routeviews.org/> on Oct 5, 2010. If interested, you can find other snapshots at <http://archive.routeviews.org/>. The peering structure of the router is at <http://www.routeviews.org/peers/route-views.oregon-ix.net.txt>. You can find IP-to-AS mappings at [http://iplane.cs.washington.edu/data/ip\\_to\\_as\\_mapping.txt](http://iplane.cs.washington.edu/data/ip_to_as_mapping.txt).

- Find the AS number corresponding to UMass from the router's routing table and state how you found it.
- What is the IP address of the next-hop from the router to UMass? Is the next-hop in the same AS as the router or in a different AS?

- (c) How does the router know how to reach that next-hop address?
- (d) List all the prefixes belonging to the UMass AS. Why are these not contiguous (defined formally in part i)?
- (e) Use `traceroute` to trace the route from UMass to the router. Is the current route from UMass to the router same as the reverse route from the router to UMass?
- (f) Below is the output of a traceroute from UMass to the router done on the same day as the routing table dump. What are the last three ASes on this path? Why is this AS path different from the AS path from the router to UMass?

```

traceroute to 128.223.51.103 (128.223.51.103), 64 hops max, 40 byte packets
 1  cs-gw (128.119.240.254)  23.338 ms  25.222 ms  3.567 ms
 2  lgrc-rt-106-8-v1263.gw.umass.edu (128.119.3.154)  2.701 ms  1.812 ms  3.799 ms
 3  border4-rt-gi7-1.gw.umass.edu (128.119.2.194)  134.800 ms  233.333 ms  223.248 ms
 4  nox300gw1-v1-535-nox-umass.nox.org (192.5.89.101)  7.627 ms  6.175 ms  7.143 ms
 5  nox300gw1-peer-nox-internet2-192-5-89-222.nox.org (192.5.89.222)  212.033 ms  241.756 ms
 6  so-0-0-0.0.rtr.wash.net.internet2.edu (64.57.28.11)  39.320 ms  32.840 ms  26.275 ms
 7  so-0-0-0.0.rtr.atla.net.internet2.edu (64.57.28.6)  29.018 ms  41.725 ms  31.498 ms
 8  so-0-2-0.0.rtr.hous.net.internet2.edu (64.57.28.43)  52.401 ms  60.815 ms  57.912 ms
 9  so-3-0-0.0.rtr.losa.net.internet2.edu (64.57.28.44)  90.059 ms  114.008 ms  91.379 ms
10  v1-101.xe-0-0-0.core0-gw.pdx.oregon-gigapop.net (198.32.165.65)  135.097 ms  109.899 ms
11  v1-105.uonet9-gw.eug.oregon-gigapop.net (198.32.165.92)  108.536 ms  109.971 ms  141.746 ms
12  ge-5-2.uonet1-gw.uoregon.edu (128.223.3.1)  111.904 ms  189.013 ms  135.633 ms
13  ge-0-2.route-views.uoregon.edu (128.223.51.103)  111.791 ms * 136.842 ms

```

- (g) What is the best route to UMass among the routes in the table? Why do you think this route was selected by the router?
- (h) Using the routing table dump and `dig` or `nslookup`, infer which ISPs provide Internet service to UMass.
- (i) We call a set of prefixes  $P_i = a_i.b_i.c_i.d_i/m_i$  for  $0 \leq i \leq k$  as contiguous if the union of the addresses in the prefixes can be written as a prefix  $P = e.f.g.h/n$  where  $n \geq \max(m_1, \dots, m_k)$ . Find a set of two or more contiguous prefixes that have the same AS path in the routing table. Why are these prefixes being advertised separately?

3. **(Inferring AS relationships):** Use the routing table dump above to answer the following questions. Refer to the paper “*On Inferring Autonomous System Relationships in the Internet*”, L. Gao, IEEE Transactions on Networking, Vol 9, No. 6, 2001. You should be able to answer the following questions using just command-line `awk` or a simple scripting language like `perl` or `python`.

- (a) Plot the CCDF (Complementary Cumulative Distribution Function) of the AS degree, i.e., plot the fraction of ASes that have a degree  $\geq n$  for all  $n > 0$ . List the ten ASes with the highest degree and the value of their degrees. Remember to (1) exclude edges from an AS to itself, and (2) include all paths, not just the best path for each prefix.
- (b) Your task is to infer the relationship between adjacent ASes in the paths listed below in the following two steps. First, for each AS path  $ABCD$  note the transit relationships as follows. If  $C$  is the AS with the highest degree, then infer the following transit relationships:  $A \rightarrow B$ ,  $B \rightarrow C$ , and  $D \rightarrow C$ .  
Second, some pairs of ASes may have a commutative transit relationship, i.e.,  $A$  transits  $B$  and  $B$  transits  $A$ . These are called *sibling* relationships. If the transit relationship

$X \rightarrow Y$  is unilateral, then we call  $X$  a customer of  $Y$ . For all adjacent AS pairs in the paths below, infer whether they have a provider-customer or a sibling relationship.

- i. 3277 3267 174 7018
- ii. 12956 1239 3561 13909
- iii. 286 3549 1239 17033

Note that the above simplistic algorithm does not attempt to disambiguate peering relationships from provider-customer and sibling relationships.

4. (**Consistency, iBGP**): In the questions below, we define *consistency* as follows. If a router  $A$  adopts a route a destination via another router  $B$ , then  $B$  adopts the suffix of  $A$ 's route starting at  $B$  as its route to the destination. Note that a stable path assignment for the Stable Paths Problem satisfies the consistency property. Consistent routes are loop-free.

- (a) Can two routers in the same AS select different AS paths to a common prefix? If yes, give an example. If no, explain why not.
- (b) Does the consistency property disallow the condition in the previous question?
- (c) To disseminate routes inside an AS, iBGP uses one of two approaches: *full mesh*, or *route reflection*. Explain each in a couple sentences.
- (d) Consider the iBGP configuration shown in Figure 1. R1 and R2 are route reflectors and are iBGP peers of each other. C1 and C2 are route reflector clients of R1 and R2 respectively. The IGP costs of the network are marked on the links in the figure. Two routes to an external destination  $d$ , which are equally good with respect to the local preference, AS path length, and MED attributes, arrive at the routers R1 and R2. Explain why even after convergence, the routes of  $C1$  and  $C2$  are not consistent.




Figure 1: iBGP configuration using route reflectors.

- (e) Assume each AS is a single router and the MRAI timer has a value of 30s. Give an example of a scenario where an AS forwards packets along an inconsistent loopy route to a destination for 60s. Your example must not have dispute wheels and all ASes must have a (consistent) route to the destination eventually.