

4th Assignment: Network Protocols and Architectures WS 10/11

Question 1: (5 + 5 + 5 + 5 + 5 + 5 + 5 = 35 points) *Routing / The border gateway protocol*

- Explain briefly the difference between routing and forwarding. What is a router, a host and an interface? Do routers have IP addresses? If yes, how many?
- Why are policies in the area of Intra-AS-routing protocols (e.g., OSPF) of little importance?
- Why do policies on the other hand play a more important role within Inter-AS-traffic compared to global traffic optimization?
- Why doesn't it make sense to optimize routes for minimum (packet) delay in intra AS routing protocols?
- How does BGP—as a member of the path-vector protocol family—bypass the problem of routing loops?
- The de-facto standard for Inter-AS routing is BGP. Why is it difficult to enforce alternatives to BGP or switch completely to a different protocol that is incompatible with BGP?
- Why are private address spaces (10/8, 172.16/12, 192.168/16) needed, which are not routed in the Internet?

Question 2: (5 + 5 + 10 + 5 = 25 points) *Properties of BGP*

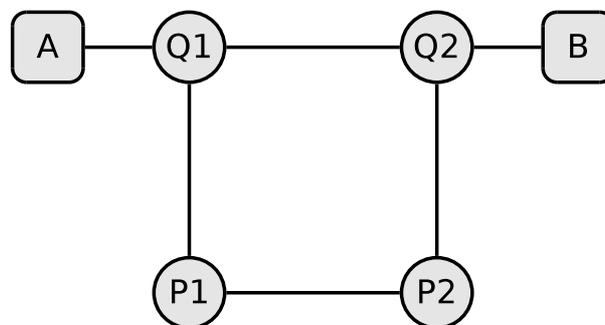


Figure 1: Intra-AS setup

Consider the topology shown in Figure 1. All edges have a weight of 5. An intra-AS routing protocol should be used.

- Identify the cost-optimal route from A to B.
- Suppose that the edge weight between Q1 and Q2 increases to 20. Given this case, state the cost optimal route from A to B. Which path will a packet actually take between A and B? Explain why.

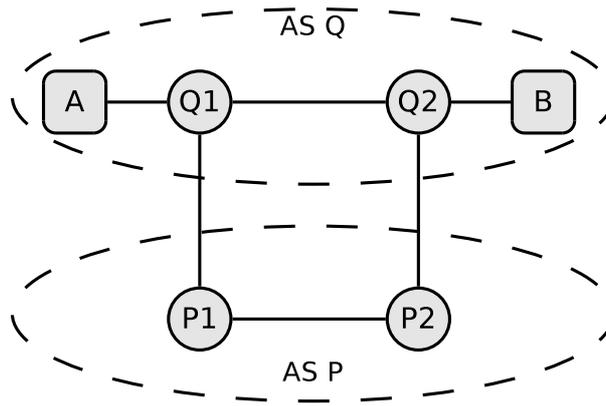


Figure 2: BGP Setup

As shown in Figure 2, the network is now split into two autonomous systems (AS). BGP is used as routing protocol between those autonomous systems.

- (c) The edge weight between Q1 and Q2 is still 20. How will a packet from A to B be routed? Explain why.
- (d) Supposing that the connection between Q1 and Q2 drops out. How are packets routed from A to B in this case? Explain why.

Question 3: (5 + 20 + 5 + 10 = 40 points) *BGP Experiment*

We are going to explore BGP in a more practical manner by visualizing the route to the Universidad de Matanzas, Cuba. First we need the IP address of the target host:

```
$ host www.umcc.cu
www.umcc.cu has address 200.14.52.71
```

Next we examine the route of a host (in this example is the source a machine in the network of TU-Berlin) to the target by using `traceroute`¹. As you can see, we get the DNS names and IP addresses of the intermediate routers, but actually we would be interested in the AS numbers. More coming next.

```
$ traceroute 200.14.52.71
traceroute to 200.14.52.71 (200.14.52.71), 30 hops max, 40 byte packets
 1 firebird.net.t-labs.tu-berlin.de (130.149.220.126) 0.231 ms 0.180 ms 0.163 ms
 2 130.149.235.1 (130.149.235.1) 0.680 ms 0.783 ms 0.788 ms
 3 xr-tub1-ge8-3.x-win.dfn.de (188.1.33.81) 0.758 ms 0.986 ms 1.215 ms
 4 zr-pot1-te0-7-0-3.x-win.dfn.de (188.1.144.221) 64.928 ms 64.823 ms 64.806 ms
 5 decix.mpd01.fra03.atlas.cogentco.com (80.81.192.63) 28.294 ms 28.399 ms 27.906 ms
 6 te0-1-0-7.ccr21.fra03.atlas.cogentco.com (130.117.1.54) 15.010 ms 14.831 ms 14.446 ms
 7 te0-2-0-6.ccr21.dca01.atlas.cogentco.com (154.54.31.237) 114.253 ms 114.051 ms 114.255 ms
 8 te3-1.ccr01.atl01.atlas.cogentco.com (154.54.28.202) 127.165 ms \
   te9-8.ccr01.atl01.atlas.cogentco.com (154.54.42.194) 123.747 ms \
   te8-1.ccr01.atl01.atlas.cogentco.com (154.54.2.45) 124.968 ms
 9 te3-1.ccr01.mia01.atlas.cogentco.com (154.54.30.34) 139.193 ms \
   te3-4.ccr01.mia01.atlas.cogentco.com (154.54.24.162) 140.966 ms \
   te3-1.ccr01.mia01.atlas.cogentco.com (154.54.30.34) 138.971 ms
10 vl3528.na01.b019263-0.mia01.atlas.cogentco.com (38.20.36.106) 137.775 ms 139.579 ms 140.049 ms
11 38.100.206.106 (38.100.206.106) 142.021 ms 142.636 ms 142.607 ms
12 newcom1.newcom-intl.com (204.14.40.1) 143.087 ms 142.220 ms 142.187 ms
13 mia-static-204-14-41-196.newcom-intl.com (204.14.41.196) 717.459 ms 717.577 ms 717.426 ms
14 200.0.16.102 (200.0.16.102) 718.917 ms 690.307 ms 692.275 ms
15 200.0.16.113 (200.0.16.113) 689.046 ms 689.037 ms 693.635 ms
16 200.0.16.141 (200.0.16.141) 692.127 ms 696.991 ms 692.355 ms
17 * * *
18 * * *
19 * * *
20 * * *
21 200.14.48.129 (200.14.48.129) 961.508 ms 1836.570 ms 1890.243 ms
22 * * *
```

¹<http://en.wikipedia.org/wiki/Traceroute>

Connect now by using *telnet* to *route-server.belwue.de* Hint: Belwue network is the national university network of Baden-Wuerttemberg. The above stated server provides you with an emulation of a Cisco IOS shell with the possibility to explore BGP routes from there to any IP addresses in the world. Enter the following command at the prompt:

```
route-server>show ip bgp 200.14.52.71
```

- (a) Which AS are hit on the BGP route to *www.umcc.cu*? Copy the result to your solution.
- Hint: The line that starts with 680 represents the BGP route. AS 680 is the number of German Research Network (DFN), through which traffic would pass first, if it has left Belwue network. Behind it, the numbers of the other AS on that BPG route are listed. To get more detailed information about them, please visit <http://www.arin.net/> (US), or <http://www.ripe.net/whois> (Europe). The syntax to query the RIPE database to get information about the AS number needs a capital “AS” as prefix.
- (b) Visualize the result of traceroute in a drawing with each AS as a dashed ellipses, routers as small circles and links as lines. Include the IP addresses and AS numbers and the location information that you may guess from the routers name. Hint:
- `route-server> traceroute 200.14.52.71`
 - “fra, mia, atl, ...” are abbreviations representing city names. “fra” is equivalent to Frankfurt/Main. Hint: Abbreviations for city names in DNS names of routers are often chosen according to airport codes².
 - Instead of using the complete IP address scheme, you can also use a prefix i.e. 188.1.0.0/16 and label the router with the remaining part of the IP address, e.g. 33.81, 144.221 and 145.137.
 - You can use geolocation databases like http://www.maxmind.com/app/locate_ip in addition (those information might not be always accurate!).
- (c) Compare the route established from *route-server.belwue.de* to *www.umcc.cu* with the above route specified by the TU-Berlin towards *www.umcc.cu* computer. Which part of the paths are identical? Where do they differ? Finally quit the telnet session with:
- ```
route-server> exit
```
- (d) Find one curiosity in the traceroute provided on this assignment sheet and explain it. Hint: you should not explain the output of traceroute but rather point to oddities that one would not expect. Examples could be: routing loops, “special” links that were traversed, unexpected latencies, slow processing / congestion, routing over Australia, ... Please note that we did not modify the output of traceroute other than inserting a \to indicate an artificial line-break. The output does not contain any syntactical errors.

**Due Date: Thursday, November, 25th 2010 only until 13:55 h s. t.**

- **As PDF files (no MS Office or OpenOffice files):** Uploaded via ISIS (<https://www.isis.tu-berlin.de/course/view.php?id=3584>)
- **On paper:** Postbox in the Telefunkenhochhaus (basement, behind the doorman right)
- Put your name, StudentID number (Matrikelnummer) **and** the name of your tutor on your solution.

---

<sup>2</sup>[http://www.johncletheroe.org/usa\\_can/airports/codes.htm#us](http://www.johncletheroe.org/usa_can/airports/codes.htm#us)