TECHNISCHE UNIVERSITÄT BERLIN
Fakultät IV – Elektrotechnik und Informatik
Fachgebiet Intelligente Netze
Prof. Anja Feldmann, Marco Canini, Petr Kuznetsov, Stefan Schmid
Arne Ludwig, Oliver Hohlfeld, Juhoon Kim, Thomas Krenc, Enric Pujol, Felix Poloczek

7th Assignment: Network Protocols and Architectures, WS 12/13

Question 1: $(5+5+5+5+5 = 25$ points) *Routing / The border gateway protocol*

(a) Why are policies in the area of intra-AS routing protocols (e. g., OSPF) of little importance?

(b) Why do policies on the other hand play a more important role within Inter-AS traffic compared to global traffic optimization?

(c) Why doesn't it make sense to optimize routes for minimum (packet) delay in intra-AS routing protocols?

(d) How does BGP—as a member of the path-vector protocol family—bypass the problem of routing loops?

(e) 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?

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

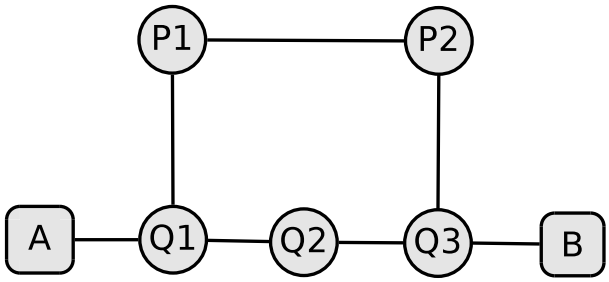

Figure 1: Intra-AS setup

Consider the topology shown in Figure 1. All edges have a weight of 2. An intra-AS routing protocol should be used.

(a) Identify the cost-optimal route from A to B.

(b) Suppose that the edge weight between Q2 and Q3 increases to 6. 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.

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 Q2 and Q3 is still 6. How will a packet from A to B be routed? Explain why.

(d) Supposing that the connection between Q2 and Q3 drops out, how are packets routed from A to B in this case? Explain why.

Question 3: $(5+20+10+5+10 = 50$ points) *BGP experiment*

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

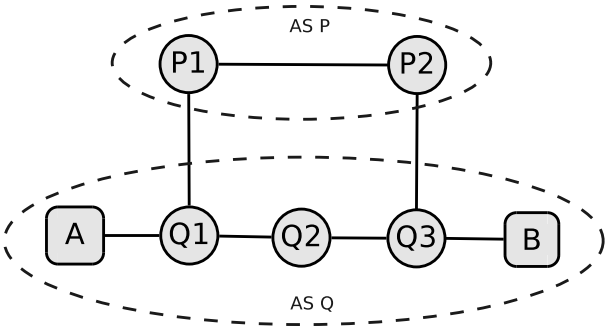

Figure 2: BGP setup

```
$ host www.udg.co.cu
www.udg.co.cu has address 200.14.53.5
```

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 200.14.53.5
 traceroute to 200.14.53.5 (200.14.53.5), 30 hops max, 60 byte packets
 1  firebird.net.t-labs.tu-berlin.de (130.149.220.126)  0.248 ms  0.171 ms  0.570 ms
 2  130.149.235.1 (130.149.235.1)  9.493 ms  9.498 ms  9.766 ms
 3  xr-tub2-te2-4.x-win.dfn.de (188.1.235.117)  2.602 ms  2.506 ms  2.490 ms
 4  cr-tub1-te0-7-0-0.x-win.dfn.de (188.1.144.165)  2.474 ms  2.458 ms  2.442 ms
 5  cr-fra1-te0-7-0-2.x-win.dfn.de (188.1.145.222)  10.398 ms cr-fra1-te0-7-0-1.x-win.dfn.de (188.1.146.2)  10.384 ms cr
 6  ae53.edge5.Frankfurt1.Level3.net (212.162.4.5)  10.267 ms  10.300 ms  10.228 ms
 7  vlan90.csw4.Frankfurt1.Level3.net (4.69.154.254)  10.207 ms  11.363 ms vlan80.csw3.Frankfurt1.Level3.net (4.69.154.1
 8  ae-71-71.ebr1.Frankfurt1.Level3.net (4.69.140.5)  10.417 ms  10.402 ms  10.387 ms
 9  ae-46-46.ebr2.Paris1.Level3.net (4.69.143.138)  19.096 ms ae-47-47.ebr2.Paris1.Level3.net (4.69.143.142)  19.148 ms
10  ae-44-44.ebr2.Washington1.Level3.net (4.69.137.62)  99.027 ms ae-42-42.ebr2.Washington1.Level3.net (4.69.137.54)  98
11  ae-82-82.csw3.Washington1.Level3.net (4.69.134.154)  97.432 ms ae-72-72.csw2.Washington1.Level3.net (4.69.134.150)
12  ae-42-90.car2.Washington1.Level3.net (4.69.149.196)  98.779 ms ae-22-70.car2.Washington1.Level3.net (4.69.149.68)  9
13  INTELSAT-IN.car2.Washington1.Level3.net (4.79.17.54)  100.908 ms  98.684 ms  100.040 ms
14  mtn-rt0003_ge-8-1.intelsatone.net (209.159.170.189)  101.573 ms  100.336 ms  101.555 ms
15  mtn-rt0003_ge-8-1.intelsatone.net (209.159.170.189)  99.018 ms  100.371 ms  100.350 ms
16  80.255.62.38 (80.255.62.38)  627.661 ms  629.508 ms  627.411 ms
17  200.0.16.82 (200.0.16.82)  628.766 ms  628.398 ms  629.713 ms
18  200.0.16.141 (200.0.16.141)  645.435 ms  645.846 ms  645.358 ms
```

Connect now by using *telnet* to *route-server.ip.tiscali.net* 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 address in the world. Enter the following command at the prompt:

```
route-server.as3257.net>show ip bgp 200.14.53.5
```

(a) Which ASes are hit on the BGP route to www.udg.co.cu? State the AS number and the corresponding name of the AS. Copy the result of `show ip bgp` to your solution.

- Hint: The line that starts with 3257 represents the BGP route. AS 3257 represents the Tiscali backbone network. To receive more detailed information about AS numbers, 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.

[1] `http://en.wikipedia.org/wiki/Traceroute`

(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.

Hints:

- `route-server.as3257.net > traceroute 200.14.53.5`
- The AS numbers are given in the traceroute output. In general, the IP to AS number mapping can be obtained by invoking `show ip bgp` for each IP address or by using a web service[2]. For a quick verification, any of the IP addresses out of the network of TU Berlin (e. g., 130.149.7.201) should be mapped to ASN 680, which is DFN.
- "fra", "Paris", ... represent city names. "fra" is equivalent to Frankfurt/Main. Hint: Abbreviations for city names in DNS names of routers are often chosen according to airport codes[3].
- Instead of using the complete IP address scheme, you can also use a prefix, e. g., 188.1.0.0/16 and label the router with the remainig part of the IP address, e. g., 33.81, 144.221, and 145.137.
- You can use geolocation databases like `http://www.maxmind.com/` or `http://www.iplocation.net/` in addition. (Those information might not be always accurate!)

(c) Compare the route established from route-server.ip.tiscali.net to www.udg.co.cu with the above route specified by the TU Berlin towards www.udg.co.cu computer. Which part of the paths are identical? Where do they differ?

(d) To which AS does the www.udg.co.cu host belong to?

(e) You will notice a significant RTT increase in both traceroute outputs. Speculate on the reason for this increase in RTT.

Due Date: Thursday, December, 13th 2012 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=7028`)
- **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.

[2] `http://asn.cymru.com/cgi-bin/whois.cgi`
[3] `http://en.wikipedia.org/wiki/List_of_airports_by_IATA_code`