

CN 2019 Problem Sheet #3

Problem 3.1: basic IPv6 network in mininet

(4+2+2+1+1 = 10 points)

Mininet by default automatically assigns IPv4 addresses to hosts and establishes connectivity between the hosts by emulating a local area network. (For the curious students, Mininet by default uses OpenFlow switches with an external controller.)

Creating an IPv6 network with the current version of Mininet requires to manually configure network interfaces and forwarding rules. The Linux command line tool of choice to configure the IP networking stack is called `ip` (you will often find explanations online how to use tools like `ifconfig` or `route` but we do not use them).

a) Write a Mininet script that establishes the following topology:

```
2001:638:709:a::/64 2001:638:709:f::/64 2001:638:709:b::/64
.-----|.-----|.-----|.
h1 ----- s1 ----- r1 ----- s0 ----- r2 ----- s2 ----- h2
eth0          eth0 eth1          eth0 eth1          eth0
```


Host `h1` is connected via switch `s1` to router `r1` and host `h2` is connected via switch `s2` to router `r2`. The routers `r1` and `r2` are connected via switch `s0`. The IPv6 networks use the prefixes shown above. Assign the following IPv6 addresses and prefixes to the hosts and routers:

```
h1-eth0 2001:638:709:a::1/64
r1-eth0 2001:638:709:a::f/64
r1-eth1 2001:638:709:f::1/64
r2-eth0 2001:638:709:f::2/64
r2-eth1 2001:638:709:b::f/64
h2-eth0 2001:638:709:b::1/64
```

Use the `ip(8)` command to configure the IPv6 addresses and the IPv6 forwarding table entries that are necessary to establish connectivity. Make sure that the routers have forwarding enabled (`sysctl -w net.ipv6.conf.all.forwarding=1`). It is recommended to use the 'Mid-level API' of Mininet.

Send the configuration commands via `h1.cmd()` etc. to the hosts before you call `net.start()`. Verify that your IPv6 network works by running `ping` and `traceroute` (or `mtr`) from all nodes to each other. Make sure you use IPv6 and not accidentally IPv4. (You can disable the automatic assignment of IPv4 addresses by passing `ip=None` to `addHost()`.)

b) Add a second network 2001:638:709:e::/64 consisting of r3, s3, r4:



Use the following IPv6 addresses:

```

r3-eth0  2001:638:709:a::e/64
r3-eth1  2001:638:709:e::1/64
r4-eth0  2001:638:709:e::2/64
r4-eth1  2001:638:709:b::e/64

```

Configure the forwarding rules such that traffic from h1 flows via r1 and r2 to h2 while traffic from h2 flows via r4 and r3 to h1. Show that the traffic between h1 and h2 is using an asymmetric path.

- c) Configure the link between r1 and r2 to have an MTU of 1400 bytes. Run ping from h1 to h2 and set the packet size so that it is bigger than the MTU on the link r1 and r2 but smaller than the default MTU. Capture the packets on the interface h1-eth0. What do you observe?
- d) Run tracepath from h1 to h2. How does ping or tracepath obtain the path MTU information?
- e) Is path MTU discovery done everytime you run ping or tracepath? If not, where does the IP implementation of h1 remember the path MTU?

Here is a Mininet script that you may use as a starting point. It also includes some hints how to debug your script if things do not work out on first try.

```
1  #!/usr/bin/env python
2  """p3-template.py:
3
4  This mininet script creates a linear routed multi-hop topology:
5
6  2001:638:709:a::/64 2001:638:709:f::/64 2001:638:709:b::/64
7  -----
8
9  h1 ----- s1 ----- r1 ----- s0 ----- r2 ----- s2 ----- h2
10 eth0          eth0 eth1          eth0 eth1          eth0
11
12 Assign the following IPv6 addresses and prefixes:
13
14 h1-eth0 2001:638:709:a::1/64
15 r1-eth0 2001:638:709:a::f/64
16 r1-eth1 2001:638:709:f::1/64
17 r2-eth0 2001:638:709:f::2/64
18 r2-eth1 2001:638:709:b::f/64
19 h2-eth0 2001:638:709:b::1/64
20
21 Debugging hints:
22
23 - If you get a 'network unreachable' error, then most likely a
24 forwarding table entry is missing somewhere or wrong.
25 - If you get a 'destination unreachable' error, then most likely
26 neighbor discovery failed somewhere.
27 - Use the 'ip -6 address', 'ip -6 route', and 'ip -6 neigh'
28 commands to inspect the forwarding tables and the neighbor
29 mapping tables.
30 - If pings from h1 to h2 do not work, try to ping r1 from h1 and
31 try to ping r1 from h2. (Test whether the links work first.)
32 - You can print the result of the configuration commands to see
33 whether any errors occurred.
34 - You can run tcpdump (or even wireshark) to capture packets in
35 order to see what is going on.
36 """
37
38 from mininet.cli import CLI
39 from mininet.net import Mininet
40 from mininet.nodelib import LinuxBridge
41 from mininet.log import setLogLevel
42
43 if __name__ == '__main__':
44     setLogLevel('info')
45
46 net = Mininet(switch=LinuxBridge, controller=None)
47
48 h1 = net.addHost('h1', ip=None)
49 # create rest of the topology here
50
51 # configure IPv6 addresses and forwarding table entries here
52 print h1.cmd("ip -V")
53
54 net.start()
55 CLI(net)
56 net.stop()
```