Modeling of OSPFv3 and BGPv4 Routing Protocols

Authors: Lukáš Galbička, Adrián Novák

xgalbi01@stud.fit.vutbr.cz, xnovak1j@stud.fit.vutbr.cz

Supervisor: Ing. Vladimír Veselý, Ph.D.

Fakulta informačních technologií, VUT v Brně

ID 3

Excel@FIT 2019

ABOUT

This paper deals with modeling and simulation of OSPFv3 and BGPv4 protocols. OSPFv3 and BGPv4 are widely used routing protocols. In their newest version are treated as modern multi-address family protocols, which means they supports both IPv4 and IPv6 routing. The resulting model may be used to demonstrate routing mechanisms in rel networks. They are both implemented in OMNeT++ Discrete Event Simulator as a part of ANSA and INET frameworks. A contribution of this work is that no working model of OSPFv3 has been yet implemented in any other simulators that are similar to OMNeT++. BGPv4 is implemented in INET4 for IPv4 network layer protocol support only and there are some issues with the current version. The version of BGPv4 does not support multi address-family routing.

Topology for BGPv4 testing

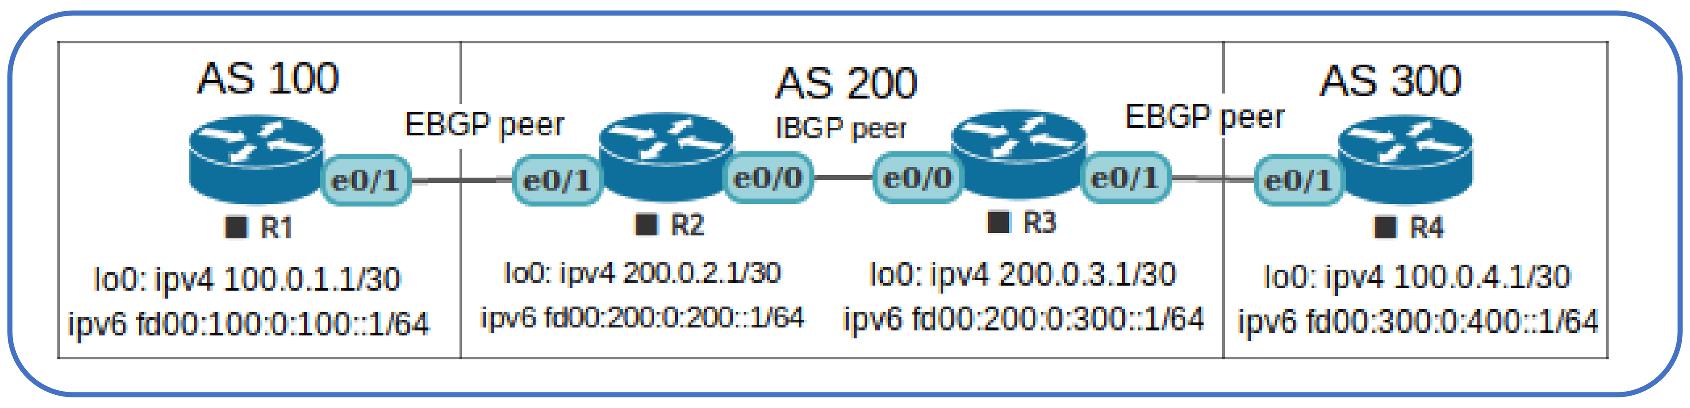

IMPLEMENTATION

Both models are implemented in OMNeT++, which is based on C++ programming language. OMNeT++ is discrete component-based simulator. Each component is implemented in C++ and models are created from these components defined in NED high-level language.

The most important contribution of both models is multi address-family routing. That means models are now able to work with IPv4 and IPv6 network layer protocol at the same time. Another new feature for BGP is the new configuration file, which allows configuring devices more comfortable than before and brings more Cisco like configuration.

BGP configuration file

```
<BGPConfig>
    <TimerParams>
        <connectRetryTime> 120 </connectRetryTime>
        <holdTime> 180 </holdTime>
        <keepAliveTime> 60 </keepAliveTime>
        <startDelay> 15 </startDelay>
    </TimerParams>
    <Devices>
        <Router name="R2" id="200.0.2.1">
            <Interfaces>
                <Interface id="eth1">
                    <Ipv4 address="10.0.20.1" netmask="255.255.255.252" />
                    <Ipv6 address="fd00:20:20::0/127" />
                </Interface>
            </Interfaces>
            <Bgp as="200">
                <Address-family id="Ipv4">
                    <Neighbor address="10.0.20.2" remote-as="200" />
                    <Network address="200.0.2.0" />
                </Address-family>
                <Address-family id="Ipv6">
                    <Neighbor address="fd00:20:20::1" remote-as="200" />
                    <Network address="fd00:200:0:300::" />
                </Address-family>
            </Bgp>
            <Route destination='200.0.3.0' netmask='255.255.255.252' interface='eth1'/>
            <Route6 destination='fd00:200:0:300::/64' interface='eth1' nexthop="fd00:20:20::1"/>
        </Router>
    </Devices>
</BGPConfig>
```

BGP real topology output

```
     Network          Next Hop            Metric LocPrf Weight Path
 *>  100.0.1.0/30     0.0.0.0                  0         32768 i
 *>  100.0.4.0/30     10.0.12.2                              0 200 300 i
 *>  200.0.2.0/30     10.0.12.2                0             0 200 i
 *>  200.0.3.0/30     10.0.12.2                0             0 200 i
R1#
```

```
     Network          Next Hop            Metric LocPrf Weight Path
 *>  FD00:100:0:100::/64
                      ::                       0         32768 i
 *>  FD00:200:0:200::/64
                      FD00:12:12::1            0             0 200 i
 *>  FD00:200:0:300::/64
                      FD00:12:12::1            0             0 200 i
 *>  FD00:300:0:400::/64
                      FD00:12:12::1                          0 200 300 i
R1#
```

BGP OMNeT++ simulation

```
BGPTestLoopback.R1.bgp._BGPRoutingTable (vector<RoutingTableEntry *>) size=4
  elements[4] (inet::bgp::RoutingTableEntry *)
    [0] BGP - Destination: 100.0.1.0/255.255.255.252 , PathType: IGP , NextHops: <unspec> , AS:
    [1] BGP - Destination: 200.0.2.0/255.255.255.252 , PathType: EGP , NextHops: 10.0.12.2 , AS: 200
    [2] BGP - Destination: 200.0.3.0/255.255.255.252 , PathType: EGP , NextHops: 10.0.12.2 , AS: 200
    [3] BGP - Destination: 100.0.4.0/255.255.255.252 , PathType: EGP , NextHops: 10.0.12.2 , AS: 200 300
```

```
BGPTestLoopback.R1.bgp._BGPRoutingTable6 (vector<RoutingTableEntry6 *>) size=4
  elements[4] (inet::bgp::RoutingTableEntry6 *)
    [0] BGP - Destination: fd00:100:0:100::/64 , PathType: IGP , NextHops: <unspec> , AS:
    [1] BGP - Destination: fd00:200:0:200::/64 , PathType: EGP , NextHops: fd00:12:12::1 , AS: 200
    [2] BGP - Destination: fd00:200:0:300::/64 , PathType: EGP , NextHops: fd00:12:12::1 , AS: 200
    [3] BGP - Destination: fd00:300:0:400::/64 , PathType: EGP , NextHops: fd00:12:12::1 , AS: 200 300
```

OSPFv3 real topology output

Link-State databse

```
        OSPFv3 1 address-family ipv6 (router-id 1.1.1.1)

                Router Link States (Area 1)

ADV Router       Age         Seq#        Fragment ID  Link count  Bits
 1.1.1.1          8          0x80000001  0            2           None
 2.2.2.2          8          0x80000003  0            1           B
 4.4.4.4          8          0x80000005  0            1           B

                Net Link States (Area 1)

ADV Router       Age         Seq#        Link ID    Rtr count
 2.2.2.2          8          0x80000001  3          2
 4.4.4.4          8          0x80000001  3          2

                Inter Area Prefix Link States (Area 1)

ADV Router       Age         Seq#        Prefix
 2.2.2.2         64          0x80000003  2222::/64
 2.2.2.2         8           0x80000001  3333::/64
 4.4.4.4         64          0x80000001  3333::/64
 4.4.4.4         8           0x80000003  2222::/64

                Link (Type-8) Link States (Area 1)

ADV Router       Age         Seq#        Link ID    Interface
 1.1.1.1         64          0x80000003  4          Et0/1
 4.4.4.4         64          0x80000002  3          Et0/1
 1.1.1.1         64          0x80000004  3          Et0/0
 2.2.2.2         64          0x80000003  3          Et0/0

                Intra Area Prefix Link States (Area 1)

ADV Router       Age         Seq#        Link ID    Ref-lstype  Ref-LSID
 2.2.2.2         64          0x80000001  3072       0x2002      3
 4.4.4.4         64          0x80000001  3072       0x2002      3
```

Routing table

```
1111::/64 [0/0]
 via Ethernet0/0, directly connected
1111::1/128 [0/0]
 via Ethernet0/0, receive
2222::/64 [110/20]
 via FE80::A8BB:CCFF:FE00:200, Ethernet0/0
3333::/64 [110/20]
 via FE80::A8BB:CCFF:FE00:400, Ethernet0/1
4444::/64 [0/0]
 via Ethernet0/1, directly connected
4444::1/128 [0/0]
 via Ethernet0/1, receive
FF00::/8 [0/0]
 via Null0, receive
```

OSPFv3 OMNeT++ simulation

Link-State databse

```
Router Link States (Area 0.0.0.1)
ADV Router   Age   Seq#         Fragment ID    Link count    Bits
1.1.1.1      9     0x80000002   0              1             None
2.2.2.2      9     0x80000001   0              1             B
4.4.4.4      9     0x80000001   0              1             B

Net Link States (Area 0.0.0.1)
ADV Router   Age   Seq#         Link State ID  Rtr count
2.2.2.2      9     0x80000001   0.0.0.101      2
4.4.4.4      9     0x80000001   0.0.0.101      2

Inter Area Prefix Link States (Area 0.0.0.1)
ADV Router   Age   Seq#         Prefix
2.2.2.2      49    0x80000001   2222::/64
4.4.4.4      49    0x80000001   3333::/64
4.4.4.4      9     0x80000003   2222::/64
2.2.2.2      9     0x80000003   3333::/64

Link (Type-8) Link States (Area 0.0.0.1)
ADV Router   Age   Seq#         Link State ID  Interface
1.1.1.1      49    0x80000001   0.0.0.101      eth0
2.2.2.2      49    0x80000001   0.0.0.101      eth0
1.1.1.1      49    0x80000001   0.0.0.102      eth1
4.4.4.4      49    0x80000001   0.0.0.101      eth1

Intra Area Prefix Link States (Area0.0.0.1)
ADV Router   Age   Seq#         Link ID        Ref-lstype    Ref-LSID
2.2.2.2      50    0x80000001   0.0.0.2        0x2002        0.0.0.101

4.4.4.4      50    0x80000001   0.0.0.2        0x2002        0.0.0.101
```

Routing table

```
routeList (IPv6Route *>)
  elements[6] (inet::IPv6Route *)
    [0] = C 1111::/64 is directly connected, eth0
    [1] = C 4444::/64 is directly connected, eth1
    [2] = O 2222::/64 [255/2] via fe80::a8bb:ccff:fe00:200, eth0
    [3] = O 3333::/64 [255/2] via fe80::a8bb:ccff:fe00:400, eth1
    [4] = S fe80::/10 is directly connected, eth0
    [5] = S fe80::/10 is directly connected, eth1
```

Topology for OSPFv3 testing

AREA 0 : R2, R3, R4

AREA 1 : R1, R2, R4

Area Border Routers: R2, R4

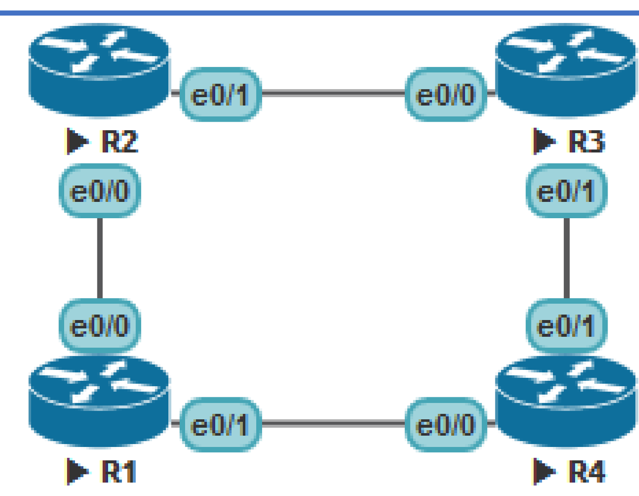

TESTING

For testing of OSPFv3 module, we used real topology consisting of four Cisco routers R1-R4 divided into two areas. The same topology is built in OMNeT++ where simulation of topology takes place. Two tested features which show the right functionality of simulated OSPFv3 protocol is correctly built Link-state database and computed routing table.

We used two methods for testing of BGPv4 module. The first method is to compare the output from real network topology recorded by Wireshark software and output from simulation models created in OMNeT++ framework. The second method compares routing tables of converged real network topology with one of the simulation models.