

Enhanced Interior Gateway Routing Protocol for OMNeT++

Vladimír Veselý, Jan Bloudíček and Ondřej Ryšavý

Faculty of Information Technology, Brno University of Technology, Božetěchova 2, Brno, Czech Republic

Keywords: Enhanced Interior Gateway Routing Protocol, EIGRP, OMNeT++, Dynamic Routing Protocol, Routing, Simulation, ANSA.

Abstract: Cisco's EIGRP is a hybrid routing protocol between distance vector and link-state routing protocols. EIGRP offers routing based on composite metric, which takes into account multiple factors and allows more granular and precise routing decisions based on the current state of the network. Cisco released basic specification of EIGRP as IETF's RFC draft in the beginning of 2013. This paper introduces one of the first freely available EIGRP design and a simulation model implementation that can be run and tested within OMNeT++ discrete event simulator.

1 INTRODUCTION

Network layer serves the purpose of end-to-end data delivery. Routers employ routing tables to make correct routing decisions by forwarding hop-by-hop the packet closer to receiver. Dynamic routing protocols maintain up-to-date content of routing tables by exchanging updates about known networks.

Routing protocols for traditional wired networks could be divided into three categories: a) **distance-vector** where routing is based on information provided by neighbors and each route has one attribute representing distance of network from a given router; b) **path-vector** which is the same as distance vector but routes have more than one attribute; and c) **link-state** where every router maintains independent view on topology and computes the shortest path tree towards all other nodes. Additional typology of routing protocol is according to type of deployment: a) **interior gateway protocols (IGP)** for routing within one administrative domain; b) **exterior gateway protocols (EGP)** for routing between autonomous systems (AS). Among typical representants belong:

- Routing Information Protocol (RIPv2 for IPv4 (Malkin, 1998), RIPv6 for IPv6 (Malkin & Minnear, 1997)) – Distance-vector routing protocol that works with hop-count as the metric. Routes with metric 16 or more are considered unreachable;
- Babel (Chroboczek, 2011) – Babel is distance-

vector protocol specialized (but not exclusively) for wireless networks that have different metric criteria than wired networks. Metric may represent cost, number of host or any other implementation dependent route attribute. Nevertheless, routes with infinity metric 0xFFFF are considered unreachable. Babel currently supports both IPv4 and IPv6;

- Intermediate System to Intermediate System (IS-IS) (Oran, 1990) – The first link-state protocol ever, which is also capable of working with different metrics simultaneously. IS-IS was originally intended to be used with Connectionless Mode Network Service Protocol (concurrent of IP) for ISO/OSI networks, however, later was developed implementation for both IPv4 and IPv6 protocols. IS-IS is by design agnostic to used address-family and single instance can carry routing updates for various network protocols. Formerly used IS-IS metrics were delay and link errorness, current revision employs only speed of the link;
- Open Shortest Path First (OSPFv2 for IPv4 (Moy, 1998), OSPFv3 for IPv4/6 (Coltun, et al., 2008)) – OSPF started as IP alternative to IS-IS and later become industrial standard link-state routing protocol that has wide-spread deployment. OSPF uses cost as the metric, where cost is derived from the interface bandwidth. OSPF supports only IP routing updates;

- Border Gateway Protocol (BGPv4) (Rekhter & Hares, 2006) – Extends distance-vector idea by having multiple attributes accompanying the single prefix update. BGPv4 is currently the only one EGP that is being used and it is often referred as policy-control routing protocol.

Enhanced Interior Gateway Protocol (EIGRP) is the backward compatible successor of previous Cisco proprietary Interior Gateway Protocol (IGRP). It is categorized as a hybrid routing protocol which means that it is a crossover between distance-vector (topology is known based on announcement from neighbors) and link-state protocols (instead of periodic updates, topology changes are propagated immediately). Down below follows the list of main beneficial features of EIGRP:

- EIGRP employs **Diffusing Update Algorithm (DUAL)** (Garcia-Lunes-Aceves, 1993) that effectively propagates any topology change and minimizes path recomputational time;
- Currently EIGRP is the only routing protocol that guarantees loop-free topology even during the time when topology is actively converging towards a new routing state;
- EIGRP leverages its own reliable transport protocol (even for multicast data transfer);
- In the contrary to other distance-vector protocols, EIGRP is capable of sending event-driven partial bounded updates;
- It has neighbor discovery and recovery mechanism to determine route reachability via particular adjacent node;
- EIGRP contains protocol-dependent modules that allow operation over different network protocols (including IPv4 and IPv6);

The EIGRP was introduced in 1993 as a conjoint effort of Cisco and SRI International (Albrightson, et al., May 1994). Initial and later measurements revealed that it outperforms other routing protocols (i.e., speed of convergece, network bandwidth utilization, queing delay) (Xu & Trajkovic, 2011). Despite its beneficial aspects (or maybe because of them) it had been protected as one of the major Cisco intellectual properties by a bunch of patents for nearly twenty years. In the beginning of 2013, basic EIGRP design and functionality were submitted as a publicly available IETF informational draft (Savage, et al., 2013).

The project ANSA (Automated Network Simulation and Analysis) running at the Faculty of Information Technology is dedicated to develop the variety of software tools that can create simulation models based on real networks and subsequently to allow formal analysis and verification of network

configurations. This paper outlines a new simulation module, which is a part of the ANSA project and which extends functionality of the INET framework (OMNeTpp/INET, 2014) in OMNeT++ (OMNeTpp, 2014).

This paper has the following structure. The next section covers a quick overview of existing EIGRP implementations (either real or simulation ones). Section 3 deals with our contribution, mainly necessary theory, proposed design and subsequent implementation. Section 4 presents validation scenarios proving corectness of the implementation. The paper is summarized in Section 5 together with unveiling our future plans.

2 STATE OF THE ART

Currently none of vendors other than Cisco supports EIGRP in its active network devices. Despite positive campaign targeting wider EIGRP acceptance, many manufacturers and customers remain skeptical and rely on a long-time proven open solutions like OSPF or IS-IS. The one of the first publicly available open-source EIGRP routing demon is being developed at the University of Žilina (GitHub/janovic, 2013) within the scope of Quagga project (nonGNU, 2013).

A freely available demonstration tool called Easy-EIGRP (SourceForge, 2013) exists rather for educational purposes.

OPNET simulator has contained EIGRP simulation modules even before its public IETF release. However, its functionality is limited and it lacks IPv6 support for EIGRP. Nevertheless, OPNET and its simulation models were used to conduct several measurement studies comparing different routing protocols including EIGRP (Wu, 2011).

Previously described state of EIGRP deployment affirmed our decision to offer academic and enterprise community with a full-fledged EIGRP implementation with all usually employed features.

The current status of unicast routing support in OMNeT 4.4.1 and INET 2.2 framework is according to our best knowledge as follows. The IPv4 (named `networkLayer`) and IPv6 (pragmatically called `networkLayer6`) layers are already parts of INET framework. The framework contains OSPFv2 as the only available dynamic routing protocol.

During ANSA project development we have extended original simple router module to be dual-stack and enhanced it with a variety of dynamic routing protocols (RIP, RIPng, IS-IS, OSPFv3,

PIM), thus creating **ANSARouter** as the compound simulation module based on the standard behavior of Cisco routers.

The basic goal behind our effort is to support EIGRP dynamic routing protocol. Hence, we have decided to add missing functionality in form of simulation module directly connected to `networkLayer` and `networkLayer6` as depicted on Figure 1.

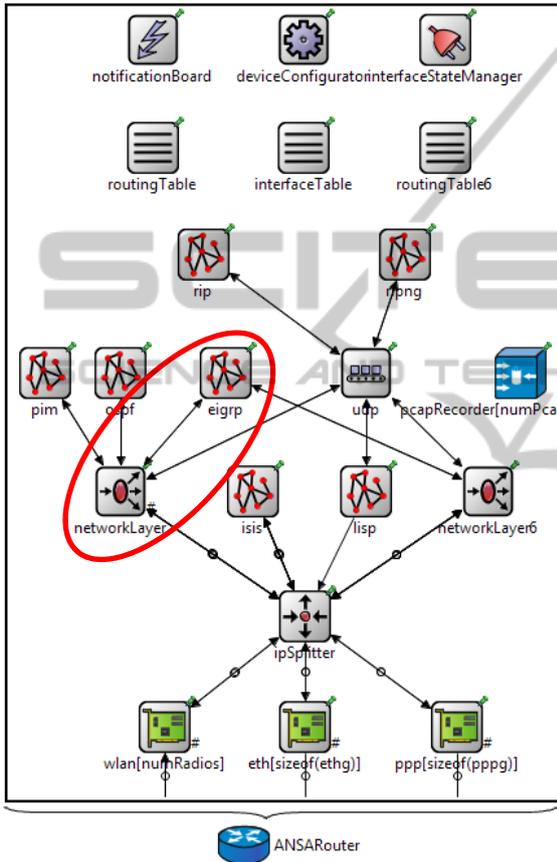


Figure 1: ANSARouter structure with highlighted contribution.

OMNeT++ state of the art prior to this paper is the result of ongoing research covered in our other articles.

3 CONTRIBUTION

We have implemented OMNeT++ compound simulation module supporting EIGRP behavior and functionality. This section provides a short theoretical background, overview of design and some implementation notes.

3.1 Theory of Operation

An EIGRP process computes a successor for every destination. A **successor** represents the next-hop router where the route to the destination via successor is loop-less and with the shortest distance. **Feasible successor (FS)** or so called backup next-hop is the router that provides loop-less route but with higher distance. To determine whether particular router is a feasible successor, the router is working with two parameters – a feasible and a reported distance. **Feasible distance (FD)** is the best known distance from a destination network to a given EIGRP router (historical minimum). **Reported distance (RD)** is distance from destination network advertised by a given EIGRP router neighbor. The router is using *FD* and *RD* to decide whether the feasible condition is satisfied or not. **Feasible condition** assumes that any route with $RD < FD$ is without any doubts loop-less. The **passive state** is the state of the destination network when the successor is known and the route is converged and usable. **Active state** is in contrast to the previous definition when the destination network does neither have a successor nor FS and the router is actively searching and computing a new successor.

The EIGRP employs composite metric which takes into account multiple route attributes. The basic composite metric consists of following four parameters: a) bandwidth (abbr. *Bw* is minimal bandwidth enroute); b) delay (abbr. *Dl* is accumulative sum of delays enroute), c) load (abbr. *Lo* is maximal traffic load in range from 1 to 255 on the links towards destination where lower is considered better), d) reliability (abbr. *Re* is minimal reliability in range from 1 to 255 on the links towards destination where higher is considered better). Parameters a) and b) are static, parameters c) and d) are dynamically recomputed every 5 minutes on certain EIGRP versions. Parameters are accompanied with K-values called weights which are unsigned byte long values where $K_4 = K_5$. Usually Cisco routers are using default composite formula for metric computation without dynamic parameters:

$$K_1 \cdot Bw + K_3 \cdot Dl$$

Complete composite formula including all parameters looks like this:

$$\left(K_1 \cdot Bw + \frac{K_2 \cdot Bw}{256 - Lo} + K_3 \cdot Dl \right) \cdot \frac{K_5}{Re + K_4}$$

The new revision of EIGRP establishes two new parameters: a) jitter (abbr. *Ji* is accumulative delay

variation enroute measured in microseconds where lower is preferred); b) energy (abbr. En is accumulative energy consumption in watts per transferred kilobit where lower is preferred). Both parameters are accompanied with K_6 weight. A new wide metric is 64 bit long in opposite to older 32 bit long standard metric and it also solves problem of standard metric when taking into account delay on links faster than 1 Gbps. Wide metric composite formula is then:

$$\left(K_1 \cdot Bw + \frac{K_2 \cdot Bw}{256 - Lo} + K_3 \cdot Dl + K_6 \cdot (En + Ji) \right) \cdot \frac{K_5}{Re + K_4}$$

When employing multicast for communication on local segment, EIGRP has either reserved address 224.0.0.10 for IPv4 or FF02::A for IPv6. EIGRP routers exchange following messages during operation:

- *EIGRP Hello* – Detects EIGRP neighbors with their settings (K-values, autonomous system number, timers and authentication) and checks their aliveness. Sent periodically every 5 seconds by default. Hold timer (period after which neighbor is considered dead) is $3 \times$ longer, and by default it is 15 seconds. Neighbor announces its own hello and hold intervals which will obey during its operation;
- *EIGRP Update* – Carries routing information that might cause receivers to start DUAL. Sent either as unicast or multicast;
- *EIGRP Ack* – Used for acknowledging *EIGRP Update*, *Query* and *Reply* messages. It is reused *EIGRP Hello* message with empty structure;
- *EIGRP Query* – If network transits to active state and router starts to search for a new successor then router starts DUAL and sends *EIGRP Queries* to neighbors usually as multicast;
- *EIGRP Reply* – This message contains the routing answer to previous *EIGRP Query*.

3.2 Design and Implementation

The EIGRP implementation works with three tables:

- **Neighbor table (NT)** – Stores information (e.g., IP address, router-id, uptime, hold-time, query count, etc.) relevant to all adjacent EIGRP routers;
- **Topology table (TT)** – The main routing information base from point of view of a given router. It contains each known network and relevant routes, their states and next-hop addresses together with their *FD* and *RD*;
- **Routing table (RT)** – A routing table is the gathering place of best routes from different

routing sources, thus the best EIGRP routes are installed here from topology table.

The compound EIGRP simulation module is divided into components depicted in Figure 2.

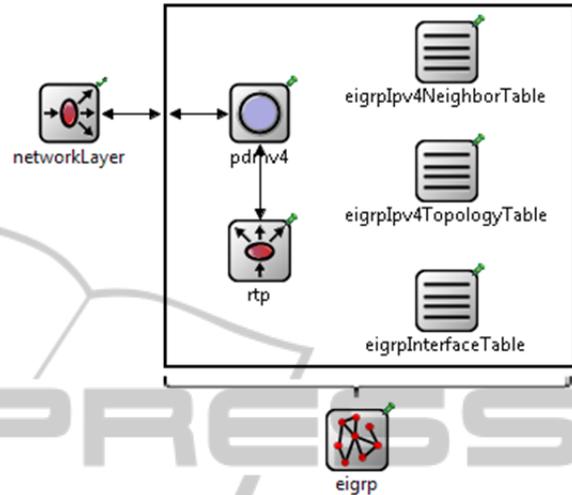


Figure 2: EIGRP simulation module structure.

A brief description of components is provided in Table 1:

Table 1: Description of EIGRP submodules.

Name	Description
pdmv4	The protocol-dependent module (PDMs) sends and receives EIGRP messages that contain IPv4 routing information. It mediates control exchange between routing table and topology table.
rtp	EIGRP uses Cisco Reliable Transport Protocol (RTP) to ensure reliable transfer of EIGRP messages. It uses sequence number and positive acknowledgement scheme to detect any gaps in transfers.
eigrpIpv4 NeighborTable	This module is emanation of neighbor table. It maintains state of all EIGRP adjacencies (i.e., neighbor IP address, state, hold timer, RTP sequence number)
eigrpIpv4 TopologyTable	EIGRP RIB which includes all learned routes, their state, <i>FDs</i> and computed successors.
eigrp InterfaceTable	Simulation module keeps settings relevant to any interface on which EIGRP is enabled (i.e., separate hello and hold timers, query count).

4 TESTING

In this section, we provide information on testing and validation of our implementation. Only two scenarios are described here really thoroughly because of limited space. Nevertheless a rich set of test scenarios is accompanied with the published source codes.

We compared results with the behavior of the referential EIGRP implementation running at Cisco routers. For this reason, we built exactly the same topology and observed (using Switched Port Analyzer and Wireshark) relevant message exchanges between real devices (Cisco 7204 as routers with IOS version c7200-adventerprisek9-mz.152-4.M2 and host stations with Windows 7 OS).

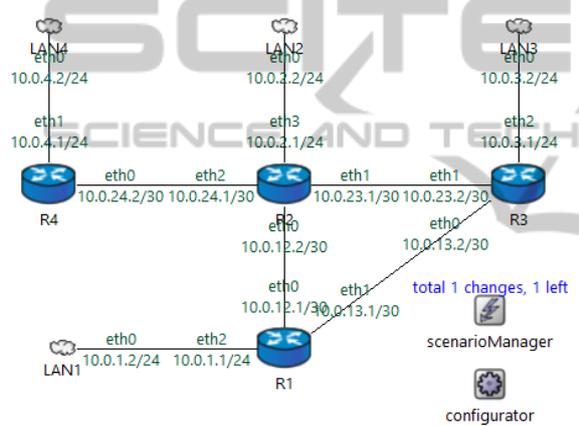


Figure 3: EIGRP testing topology.

Testing topology (see above Figure 3) consists of four EIGRP routers (marked R1, R2, R3 and R4) and four ANSAS standard hosts (LAN1, LAN2, LAN3 and LAN4) which substitutes whole separate LAN segment with dedicated IP networks.

In the first scenario, we would like to show how metric changes are being propagated. In the second scenario, we focus on topology changes.

4.1 Scenario 1: Metric Change

A typical message exchange of freshly booted routers is following:

- #1) Routers establish neighborhood by sending and receiving *EIGRP Hello* messages. Whenever a new neighbor is discovered, all relevant information is recorded and stored in NT. This fact is depicted on Figure 4 where we can observe `egrpIpv4NeighborTable` content on

router R2 just few seconds after beginning of scenario (no later than 4 seconds after the start);

- #2) Whenever neighborhood is established, routers exchange *EIGRP Updates* containing routing information to build their TTs and determine best routes towards known destinations. Reception and processing of any update is confirmed by *EIGRP Ack*. Figure 5 shows converged state of topology from the router R2's `egrpIpv4TopologyTable` point of view. Routes have known *FD*, successors and are in passive states.



Figure 4: R2's Neighbor Table prior to Scenario 1 events.

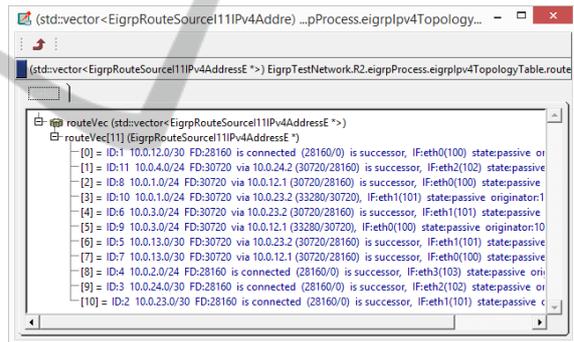


Figure 5: R2's Topology Table prior to Scenario 1 events.

We scheduled bandwidth alternation R3's eth2 interface facing LAN3 changes its *Dl* attribute from 100 to 10 000 in order to show how the change of metric influences topology (for instance content of R2's RT is depicted on Figure 6). In simulator, we use `scenarioManager` to accomplish this goal, in case of real-network, we change interface configuration.

- #3) R3 initiates DUAL, which discovers that only 10.0.3.0/24 is reachable via eth2 and propagates metric change to its neighbors R2 and R1 by sending *EIGRP Update* for network 10.0.3.0/24;
- #4) R2 acknowledges update with *EIGRP Ack*. R2's DUAL is unable to find FS, hence route transits to active state and router sends ordinary *EIGRP Query* to R1 and R4 and

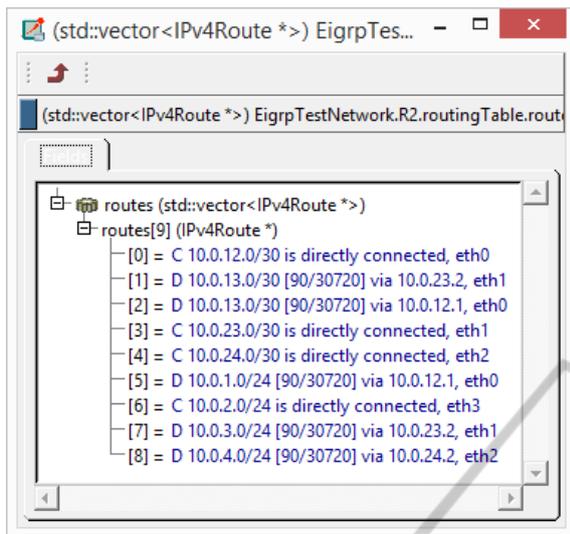


Figure 6: R2's Routing Table prior to Scenario 1 events.

poison reverse *EIGRP Query* with maximal metric towards R3. Same previous steps apply also for R1 where situation is similar – acknowledgment towards R3, DUAL marks network as active, query to R2 and poison reverse query to R3;

- #5) R1 receives *EIGRP Query* from R2 and it acknowledges it with *EIGRP Ack*. Following next, R1 responds with *EIGRP Reply* with a new metric via successor R3. Same situation repeats on R2 when replying to R1 query;
- #6) R3 receives queries from R1 and R2 and it acknowledges them. Following next, R3 finds out FS (itself) and responds with *EIGRP Replies* to R2 and R1;
- #7) R4 receives *EIGRP Query* from R2 and confirms it with *EIGRP Ack*. DUAL is unable to determine FS, thus route transits to active state. Because of split-horizon rule, there is no neighbor to query. Hence, R2 is marked as a successor due to infinity *FD*. The network transits back to passive state with a changed metric via new and old successor R2. R4 sends poison reverse *EIGRP Reply* back to R2;
- #8) R1 and R2 receive and acknowledge *EIGRP Replies* which they exchanged and store a new metric in TT;
- #9) R1 and R2 receive *EIGRP Reply* from R3 and store a new metric in TT. Because all neighbors of R1 and R2 responded to their queries, DUAL stops. Next, they both R1 and R2 update records in RTs to reflect changed metric situation of network 10.0.3.0/24.

Topology is converged and state of R2's routing table is depicted on Figure 7.

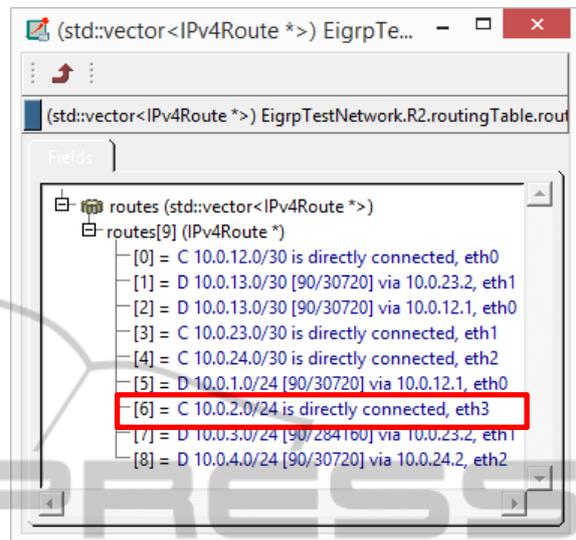


Figure 7: R2's Routing Table after Scenario 1 events.

4.2 Scenario 2: Topology Change

Scenario begins exactly same as the previous one with phase #1, when neighbors are discovered, and phase #2, when topology converges by initial routing information exchange (same content of R2's NT, TT and RT as on Figure 4, 5 and 6).

We scheduled link failure (R2's eth1) of interconnection between routers R2 and R3 for this scenario. Goal is to show how topology change is propagated from the source to other routers. Once again we accomplish this with the help of *scenarioManager* in simulator. In case of real network, we just shut down the interface. In both cases, R3's eth1 remains operational.

We have decided to omit all acknowledgement in subsequent text in order to make it clearer and easier to read. Nevertheless, all routers correctly confirm reception of *EIGRP Update*, *Query* and *Reply* messages by sending *EIGRP Ack*. Scenario continues in following manner:

- #3) Eth1 comes down on R2. *EIGRP* process goes through TT and transits all networks reachable via successor (10.0.23.0/30 and 10.0.3.0/24) on eth1 interface to active state. R2 sends *EIGRP Queries* to neighbors R1 and R4. Load balancing is enabled, thus 10.0.13.0/30 is reachable via two routes in the RT – the one that leads through R3 is removed and neighbors are notified by *EIGRP Update* messages;

- #4) R4 receives *EIGRP Query* from R2. DUAL cannot find FS for routes and because of split-horizon rule there is no other neighbor to ask. Hence, R4 sends *EIGRP Reply* stating that 10.0.23.0/30 and 10.0.3.0/24 are unreachable from its perspective;
- #5) R1 receives *EIGRP Query*. Dual finds out FS and responds back with *EIGRP Reply*. Moreover, the route to 10.0.23.0/30 via R2 is removed from RT and *EIGRP Update* about this is sent to neighbors R3 and R2. Routes on this router remain in passive state;
- #6) Integrated optimization prevents information from particular updates to be passed to DUAL. Namely previously sent *EIGRP Update* from R1 to R3, from R2 to R1, from R1 to R2 and from R2 to R4;
- #7) R2 receives *EIGRP Reply* from R4 and from R1. All replies has been received, thus routes to 10.0.23.0/24 and 10.0.3.0/30 has a new successor in R2's TT and that is R1. Those routes are propagated to R2's RT and information about change is sent to neighbors as *EIGRP Update*;
- #8) R4 receives EIGRP Update from R2 and inserts R2 as a new successor to its RT. Because of RT change, poison reverse *EIGRP Update* is sent back to R2;
- #9) Same optimization as in case of phase #6. EIRGP Updates from R2 to R1 and from R4 to R2 are omitted from DUAL processing. Content of R2's NT, TT and RT does not change for the rest of scenario and it shown on Figures 8, 9 and 10);

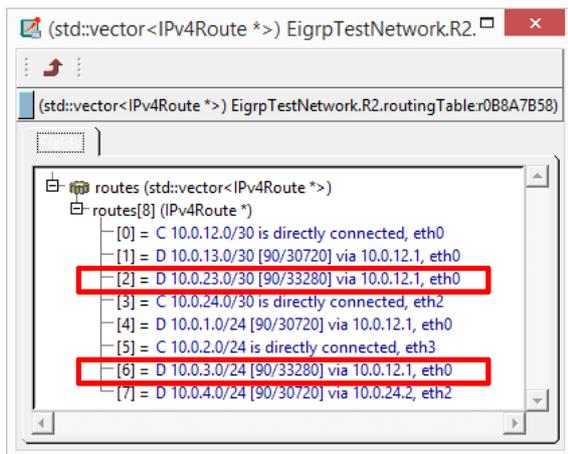


Figure 8: R2's Routing table after Scenario 2 events.



Figure 9: R2's Neighbor table after Scenario 2 events.

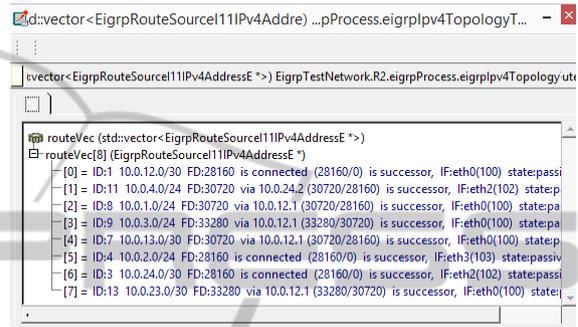


Figure 10: R2's Topology Table after Scenario 2 events.

- #10) Hold timer expires on R3, thus neighborship is terminated and R2 is removed from R3's NT. Also R3 sends goodbye *EIGRP Hello* as a preventive notification. All affected networks reachable via R2 (10.0.24.0/30, 10.0.2.0/24, 10.0.4.0/24) transit to active state and EIGRP Query is sent to remaining neighbor R1. Only exception is 10.0.12.0/30 that has another FS due to load balancing. However, its second route is removed from R3's RT and *EIGRP Update* is sent to R1;
- #11) R1 receives *EIGRP Query* and *Update* from R3. DUAL finds FS for all queried routes in R1's TT and thus no network transits to active state. *EIGRP Reply* is sent to R3 as response;
- #12) R3's DUAL collects all (single) *EIGRP Replies* (from R1). R3's TT is updated with a new successor and affected networks transit back to passive state. The best routes are introduced to R3's RT and EIGRP Update is sent to R1;
- #13) Processing of update is optimized just as in case of phase #6 and #9 on R1. Topology is converged.

4.3 Test Summary

Comparison for Scenario 1 can be observed in Table 2. Similarly description for Scenario 2 is in Table 3. Column marked with header “S → R” contains sender and receiver of a given message. In comparisons, we have focused on messages processed mostly by router R2. Nevertheless, messages that are not shown and were processed by other routers are also in correct order and without any significant deviations between simulation and real time.

The correlation of messages between simulation and real network suggests correctness of our EIGRP implementation.

Table 2: Timestamp comparison for Scenario 1.

Phase	Message	S → R	Simul. [s]	Real [s]
#3	<i>Update</i>	R3→R2	0.000	0.000
#4	<i>Query</i>	R2→R1	0.001	0.036
#5	<i>Reply</i>	R1→R2	0.002	0.068
#9	<i>Reply</i>	R4→R2	0.003	0.088

Table 3: Timestamp comparison for Scenario 2.

Phase	Message	S → R	Simul. [s]	Real [s]
#3	<i>Query</i>	R2→R1	0.000	0.040
#4	<i>Reply</i>	R4→R2	0.001	0.076
#5	<i>Reply</i>	R1→R2	0.001	0.096
#7	<i>Update</i>	R2→R1	0.002	0.116
#8	<i>Update</i>	R4→R2	0.003	0.180
#10	<i>Hello</i>	R3→R2	10.659	10.736
	<i>Query</i>	R3→R1	10.660	10.764
#11	<i>Reply</i>	R1→R3	10.661	10.812

Validation testing against the real-life topology shows just reasonable time variations (around ± 200 milliseconds). This variation observable on Cisco devices is caused by three factors: a) control-plane processing delay and internal EIGRP optimizations; b) packet pacing that guarantees constant bandwidth consumption by EIGRP process and avoids potential race conditions between EIGRP instances; and c) inaccuracy in timing of certain event in real-life network. Nevertheless, the routing outcomes of simulated and real network are exactly same when taking into account accuracy in order of seconds.

5 CONCLUSIONS

We presented an overview of the theory behind EIGRP routing protocol. The main contribution of work is a new OMNeT++ simulation module that

mimics Cisco’s EIGRP protocol implementation based on the available specification and from a reverse-engineering observations. We introduce a simulation scenario and relevant results to demonstrate its compliance with the reference Cisco IOS implementation.

EIGRP is beneficial namely for large enterprise networks because it generally consumes less resources than link-state IGPs. It is the one of the best distance-vector IGPs available and with its public release we can expect that more companies will tend to use it. For such entities, we offer polished simulation models for a reliable comparison on their network functionality which now includes also EIGRP.

We plan to carry on work towards: 1) wrap up IPv6 protocol dependent module for our EIGRP simulation module; 2) extend functionality with stub functionality, stuck-in-active support and further tune EIGRP simulation model. Additional plan is to conduct comparative evaluation of our models against those in OPNET simulator.

More information about the ANSA project is available on homepage (Brno University of Technology, 2014). All source codes including EIGRP implementation could be downloaded from GitHub repository (GitHub/kvetak, 2013).

ACKNOWLEDGEMENTS

This work was supported by the Brno University of Technology organization and by the research grants:

- FIT-S-14-2299 supported by Brno University of Technology;
- VG20102015022 supported by Ministry of the Interior of the Czech Republic;
- IT4Innovation ED1.1.00/02.0070 supported by Czech Ministry of Education Youth and Sports.

REFERENCES

- Albrightson, R., Garcia-Luna-Aceves, J. J. & Boyle, J., May 1994. EIGRP a fast routing protocol based on distance vectors. *Proceedings Network/Interop*, Volume Vol. XCIV, pp. 136-147.
- Brno University of Technology, 2014. [Online] Available at: <http://nes.fit.vutbr.cz/ansa/pmwiki.php>.
- Coltun, R., Ferguson, D., Moy, J. & Lindem, A., 2008. *RFC 5340: OSPF for IPv6*. [Online] Available at: <http://tools.ietf.org/html/rfc5340>.
- Garcia-Lunes-Aceves, J., 1993. Loop-Free Routing Using Diffusing Computations. *IEEE/ACM Transactions on Networking*, Vol. I(No. 1), pp. 130-141.

- GitHub/janovic, 2013. *janovic/Quagga-EIGRP*. [Online] Available at: <https://github.com/janovic/Quagga-EIGRP> [Accessed April 2014].
- GitHub/kvetak, 2013. [Online] Available at: <https://github.com/kvetak/ANSA> [Accessed January 2014].
- Chroboczek, J., 2011. *RFC 6126: The Babel Routing Protocol*. [Online] Available at: <http://tools.ietf.org/html/rfc6126>.
- Malkin, G., 1998. *RFC 2453: RIP Version 2*. [Online] Available at: <https://tools.ietf.org/html/rfc2453>.
- Malkin, G. & Minnear, R., 1997. *RFC 2080: RIPng for IPv6*. [Online] Available at: <https://tools.ietf.org/html/rfc2080>.
- Moy, J., 1998. *RFC 2328: OSPF Version 2*. [Online] Available at: <http://tools.ietf.org/html/rfc2328>.
- nonGNU, 2013. *Quagga Software Routing Suite*. [Online] Available at: <http://www.nongnu.org/quagga/> [Accessed April 2014].
- OMNeTpp/INET, 2014. *INET Framework | Main / Welcome to the INET Framework*. [Online] Available at: <http://inet.omnetpp.org/> [Accessed April 2014].
- OMNeTpp, 2014. *OMNeT++ Network Simulation Framework*. [Online] Available at: <http://www.omnetpp.org/> [Accessed April 2014].
- Oran, D., 1990. *RFC 1142: OSI IS-IS Intra-domain Routing Protocol*. [Online] Available at: <http://tools.ietf.org/html/rfc1142>.
- Rekhter, Y. & Hares, S., 2006. *RFC 4271: A Border Gateway Protocol 4 (BGP-4)*. [Online] Available at: <http://tools.ietf.org/html/rfc4271>.
- Savage, D. et al., 2013. *Enhanced Interior Gateway Routing Protocol*. [Online] Available at: <https://tools.ietf.org/html/draft-savage-eigrp-01>.
- SourceForge, 2013. *Easy-EIGRP | Free software downloads at SourceForge.net*. [Online] Available at: <http://sourceforge.net/projects/easy-eigrp/>
- Wu, B., 2011. *Simulation Based Performance Analyses on RIPv2, EIGRP, and OSPF Using OPNET*. [Online] Available at: http://digitalcommons.uncfsu.edu/macsc_wp/11.
- Xu, D. & Trajkovic, T., 2011. *Performance Analysis of RIP, EIGRP, and OSPF Using OPNET*. [Online] Available at: <http://summit.sfu.ca/item/10841>.

PRESS
TECHNOLOGY PUBLICATIONS