Performance Analysis of RIP, EIGRP, and OSPF using OPNET

Don Xu and Ljiljana Trajković
Simon Fraser University
Vancouver, British Columbia
Canada
E-mail: {donx, ljilja}@sfu.ca

Abstract
Routing protocols are key elements of modern communication networks. Currently deployed dynamic routing protocols that are used to propagate network topology information to the neighboring routers are *Routing Information Protocol* (RIP), *Enhanced Interior Gateway Routing Protocol* (EIGRP), and the *Open Shortest Path First* (OSPF) protocol. The choice of the right routing protocol depends on a number of parameters. In this paper, we use OPNET Modeler to analyze the performance of RIP, EIGRP, and the OSPF protocols, which are commonly deployed in Internet Protocol (IP) networks. We designed various simulation scenarios to compare their performance.

1. Introduction

Routing protocols are based on routing algorithms, which rely on various metrics to find the best path to transmit data across networks. Metrics include cost, bandwidth, maximum transmission unit (MTU), packet delay, and hop count. Routing protocols utilize a routing table to store the results of these metrics. Based whether the routing is within an A*utonomous System* (AS) or between ASs, there are two types of routing protocols: *Interior Gateway Protocols* (IGP) and *Exterior Gateway Protocol* (EGP). RIP, EIGRP, and OSPF are three commonly used IGPs. A typical EGP is the *Border Gateway Protocol* (BGP).

Routing is the process of selecting paths in a network. In packet switching networks, routing directs traffic forwarding of logically addressed packets through intermediate nodes from their source to their ultimate destination. Routing protocols are designed to select and determine the best path to each router in the network. Routers should learn the next hop to send the packets. Forwarding data should be efficient and effective. Consequently, the routing decision of a protocol is very important for network performance.

Among many routing protocols, RIP, EIGRP, and OSPF have been most widely used. In this paper, we use OPNET Modeler to simulate performance of the RIP, EIGRP, and OSPF TCP/IP Internet routing protocols. For each protocol, we designed two simulation scenarios. We collected, compared, and analyzed simulation results in the terms of network convergence, routing traffic, Ethernet delay, email upload response time, HTTP page response time, video conferencing, packet end-to-end delay, and voice packet delay.

In this paper, we compare performance of RIP, EIGRP, and OSPF using OPNET. In Section 2, we describe routing protocols. OPNET models of routing protocols are described in Section 3 while simulation scenarios are given in Section 4. Discussion of simulation results is given in Section 5. We conclude with Section 6.

2. Dynamic Routing Protocols Overview

Dynamic routing protocols play an important role in today's networks. They are used to facilitate the exchange of routing information between routers. They dynamically share information between routers, automatically update routing table when topology changes, and determine the best path to a destination. Compared to static routing, dynamic routing protocols have better scalability and adaptability and require less administrative overhead. Dynamic routing protocols allow routers to dynamically advertise and learn routes, determine available routes and identify the most efficient routes to a destination. Dynamic routing protocols have the capability to maintain the network operation in case of a failure or when network configuration or topology change [1].

"Distance vector" and "link state" are used to describe routing protocols used by routers to forward packets. There are two groups of routing protocols, based on whether the routing protocol selects the best routing path based on a distance metric (the distance) and an interface (the vector) or selects the best routing path by calculating the state of each link in a path and finding the path with the lowest total metric to the destination. Distance vector protocols evaluate the best path based on distance, which can be measured in terms of hops or a combination of metrics calculated to represent a distance value. The IP Distance vector routing protocols in use today are RIP and IGRP. In link state routing, every node constructs a map of the connectivity to the network in the form of a graph showing connectivity of the nodes to each other. Each node then independently calculates the next best logical path to every possible destination in the network [2]. The collection of best paths forms the node's routing table. Link state protocols have the routers announce their closest neighbors to every router in the network. Only a part of the table pertaining to its neighbors is distributed. EIGRP, OSPF, and *Intermediate System-Intermediate System* (IS-IS) are link state routing protocol.

2.1 Routing Information Protocol (RIP)

RIP is a distance vector dynamic routing protocol that employs the hop count as a routing metric. RIP is implemented on top of the *User Datagram Protocol* (UDP) as its transport protocol. It is assigned the reserved port number 520. RIP prevents routing loops by implementing a limit on the number of hops allowed in a path from the source to a destination. The maximum number of permitted hops is 15. Hence a hop count of 16 is considered an infinite distance. This hop number limits the size of networks that RIP may support. RIP selects paths that have the smallest hop counts. However, the path may be the slowest in the network. RIP is simple and efficient in small networks. However, it may be inefficient in larger networks. Every RIP router broadcasts to other routers the best path based on its calculation. Each router updates its own routing table by communicating with neighboring routers. RIP router transmits full updates every 30 seconds [3]. RIP may take 30–60 seconds to converge based on the features of distance vector protocols. RIP has lower power consumption and memory than some other routing protocols. RIP may be implemented in all types of routing devices. Consequently, it is a better choice in a multi-brand, mixed network.

2.2 Enhanced Interior Gateway Routing Protocol (EIGRP)

EIGRP is a Cisco proprietary routing protocol. It is based on a new route calculation algorithm called the *Diffusing Update Algorithm* (DUAL). It has features of both distance vector and link state protocols. EIGRP metrics are based on reliability, MTU, delay, load, and bandwidth. Delay and bandwidth are the basic parameters for calculating metrics [4]. EIGRP collects data from three tables. The first is the neighbors' table, which stores data about neighboring routers that are directly accessible through interfaces that are connected. The second is the topology table, which contains the aggregation of the routing tables that are gathered from all neighbors that are directly connected. It contains a list of destination networks in the EIGRP routed network and their respective metrics. The third routing table stores the actual routes to all destinations. EIGRP differs from most distance vector protocols because it does not rely on periodic route dumps. Hence, it is capable of maintaining its topology table. Information that is to be routed is only exchanged when the new neighbors adjacencies is established. The EIGRP router maintains its own routing table and tables of its neighbors [5]. The EIGRP router broadcasts to other neighbors if it cannot locate a router based on its routing database.

2.3 Open Shortest Path First (OSPF)

OSPF uses a link state routing algorithm that operates within a single AS. OSPF is an efficient IGP and may exhibit faster routing compared to RIP [6]. OSPF maintains the routing table for all connections in the network while RIP only maintains the routing table of the best path for every destination. Each OSPF router stores the local network connection state with *Link State Advertisement* (LSA) and advertises to the entire AS. Each router receives the LSA generated by all routers within the AS. The LSA collection then forms *Link State Database* (LSDB). Each LSA is the description of the surrounding network topology of a router. Hence, the LSDB reflects the AS network topology [7]. When a new router is added to the network, it will broadcast *hello* messages to every neighbor and will receive the feedback *hello* messages from its neighbors. Eventually, routers establish connections with newly added router and synchronize their routing databases. Every router broadcasts link state update messages when network topology changes. Consequently, all routers may keep same information of network topology. Every router calculates the best paths to all destinations and indicates the closet router for each transmission. OSPF is the most widely used IGP in large enterprise networks.

3. OPNET Routing Protocol Models

We used OPNET Modeler version 14.0.A for network simulations. OPNET is a comprehensive network simulation tool with a multitude of powerful functions. It enables simulation of heterogeneous networks by employing a various protocols [2]. The simulated network shown in Figure 1 consists of five subnets connected to each other with *Point to Point Protocol* (PPP) using *Digital Signal 3* (DS3, 44.736 Mb/s). Each subnet consists of Cisco 7200 routers, Cisco 3600 switches, Ethernet server, and 100BaseT LANs. These nodes are connected with Ethernet 100BaseT cables as shown in Figure 2.

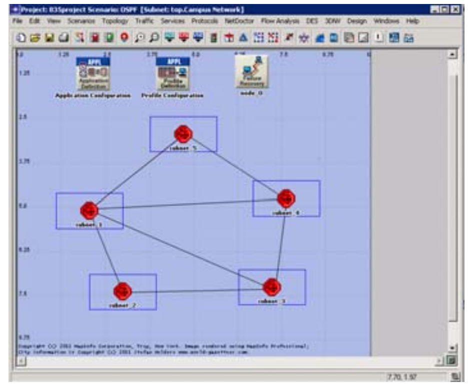

Figure 1: OPNET simulated network topology.

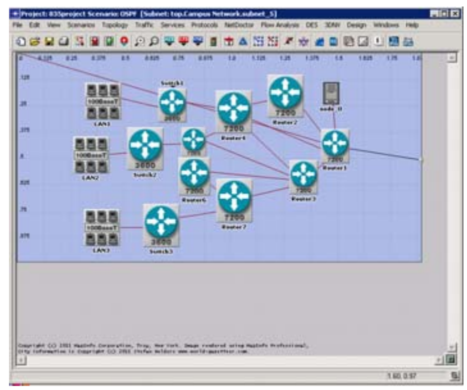

Figure 2: OPNET node model of the local area network within subnet nodes.

The first simulated network employs the RIP routing. The same model is then used to simulate EIGRP and the OSPF routing protocol. The three scenarios are: "RIP no fail", "EIGRP no fail", and "OSPF no fail". We added the failure/recovery setting (the link between *Subnet1* and *Subnet5* fails at 300 s and recovers at 500 s) to each scenario and created three additional scenarios named: RIP, EIGRP, and OSPF. The details of these six scenarios are shown in Table 1. The application configurations are shown in Table 2.

Table 1: Six OPNET simulation scenarios.

Scenario name	Routing protocol	Failure link	Fail time	Recovery time
RIP no fail	RIP	N/A	N/A	N/A
EIGRP no fail	EIGRP	N/A	N/A	N/A
OSPF no fail	OSPF	N/A	N/A	N/A
RIP	RIP	Subnet1–5	300 s	500 s
EIGRP	EIGRP	Subnet1–5	300 s	500 s
OSPF	OSPF	Subnet1–5	300 s	500 s

Table 2: OPNET application configurations.

Email	High load
HTTP	HTTP 1.1, heavy browsing
Video Conferencing	15 frames/s, 128x240 pixels
Voice	IP telephony and silence suppressed

4. Simulation Scenarios

We simulated the network convergence activity and protocol traffic using six simulation scenarios. The RIP, EIGRP, and OSPF protocol are chosen under global statistics.

Network Convergence: The EIGRP and the OSPF protocol experience the shortest and the longest network convergence times, respectively, as shown in Figure 3 and Figure 4.

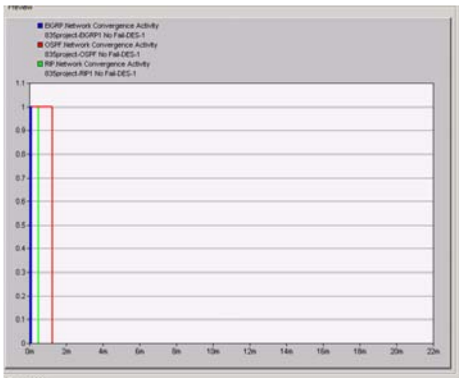

Figure 3: Network convergence activity with no failure.

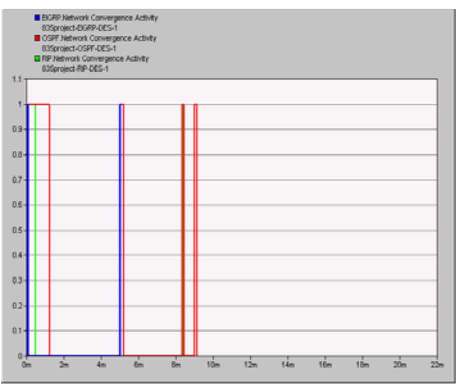

Figure 4: Network convergence activity with failure at 300 s and recovery at 500 s.

Routing Traffic: The OSPF protocol generates higher traffic compared to EIGRP and RIP. After failure/recovery occurs, the OSPF protocol generates lower traffic than EIGRP, as shown in Figure 5 and Figure 6.

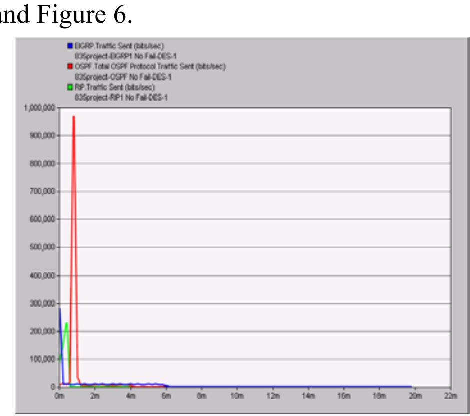

Figure 5: Protocol traffic sent (bits/s) with no failure.

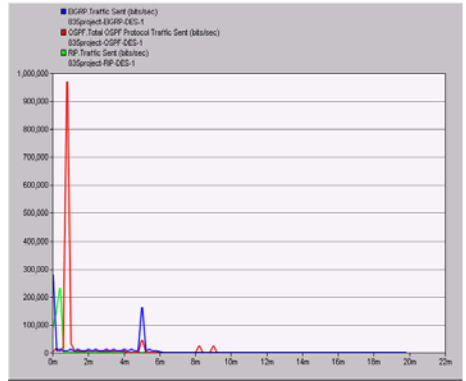

Figure 6: Protocol traffic sent (bits/s) with failure at 300 s and recovery at 500 s.

Ethernet Delay: The lowest and the highest delays are experienced by EIGRP and RIP, respectively, as shown in Figure 7.

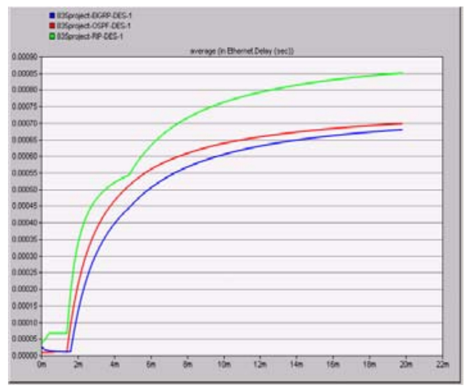

Figure 7: Average Ethernet delay (s) with failure at 300 s and recovery at 500 s.

Email Upload Response Time: The OSPF protocol exhibits the shortest response time before failure. After recovery, its response time is the highest, as shown in Figure 8.

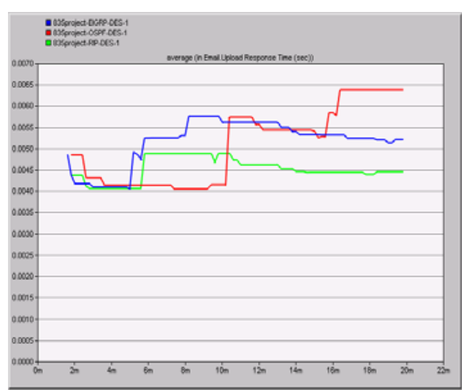

Figure 8: Average email uploads response time (s) with failure at 300 s and recovery at 500 s.

HTTP Page Response Time: The OSPF protocol exhibits the lowest response time, as shown in Figure 9.

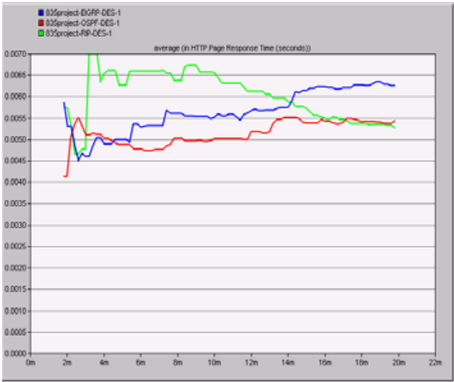

Figure 9: Average HTTP page response time (s) with failure at 300 s and recovery at 500 s.

Video Conferencing Packet End-to-End Delay: The OSPF protocol exhibits the lowest delay, as shown in Figure 10.

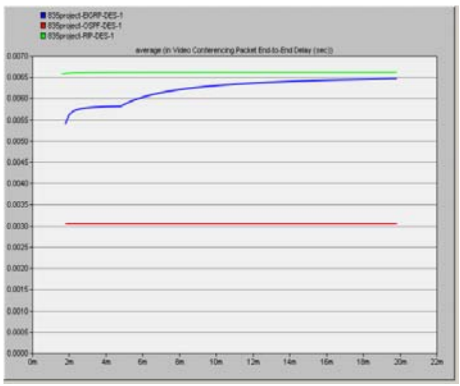

Figure 10: Average video conferencing packet end-to-end delay (s) with failure at 300 s and recovery at 500 s.

Voice Packet Delay: The RIP and the OSPF protocol experience the lowest and highest delays, respectively, as shown in Figure 11.

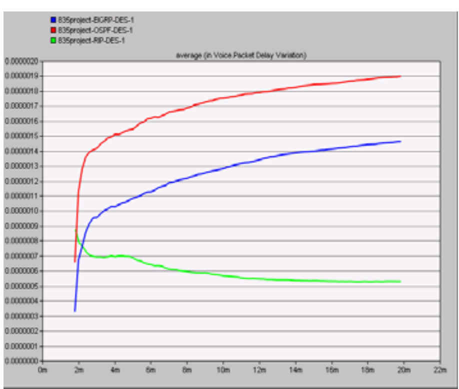

Figure 11: Average voice packet delay variation with failure at 300 s and recovery at 500 s.

5. Analysis of Simulation Results

Simulation results indicate that RIP performs better in terms of voice packet delay because it is a simple routing protocol that relies on distance vector algorithms. RIP generates less protocol traffic compared to EIGRP and OSPF, especially in medium size networks simulated in this project. RIP's weakness is slower convergence time in larger networks. This weakness may cause inconsistent routing entries and occasionally results in routing loops or metrics approaching infinity. RIP is preferred in networks smaller than 15 hops.

EIGRP performs better in terms of network convergence, routing traffic, and Ethernet delay. EIGRP has the characteristics of both distance vector and link state protocols, has improved network convergence, reduced routing protocol traffic, and less CPU and RAM utilization compared to RIP and the OSPF protocol. EIGRP has very low usage of network resources during normal operation since only *hello* packets are transmitted. When a routing table changes, its convergence time is short and it reduces bandwidth utilization. EIGRP is a Cisco proprietary protocol and, hence, a network with a non-Cisco router cannot deploy EIGRP.

OSPF performs better in terms of HTTP page response time and packet end-to-end delay for video conferencing. OSPF has large protocol overhead when updating the routing table. If there is no network change, OSPF uses very little bandwidth. OSPF is an open standard protocol and has the ability to handle large networks. Its drawback is that it relies on a more complex algorithm compared to RIP and EIGRP and requires more time to converge when building routing table and, hence, it generates additional protocol traffic. In a medium size simulated network, OSPF demands additional processing and memory requirement and consumes a large bandwidth for the initial link-state packet flooding.

6. Conclusions

In this paper, we demonstrated that OPNET Modeler can be employed by network planners to select the most suitable routing protocol for various networks and to design an optimal routing topology.

References

[1] S. G. Thornier, "Communication service provider's choice between OSPF and IS-IS dynamic routing protocols and implementation criteria using OPNET simulator," in *Proc. Second International Conference on Computer and Network Technology (*ICCNT), Bangkok, Thailand, Apr. 2010, p. 38–42.

[2] S. G. Thornier, "Dynamic routing protocol implementation decision between EIGRP, OSPF, and RIP based on technical background using OPNET Modeler," in *Proc. Second International Conference* on *Computer and Network Technology (ICCNT)*, Bangkok, Thailand, Apr. 2010, pp. 191–195.

[3] B. Fortz, J. Rexford, and M. Thorup, "Traffic engineering with traditional IP routing protocols," *IEEE Communications Magazine*, vol. 40, no. 10, pp. 118–124, Oct. 2002.

[4] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, and J. Rexford, "Netscape: traffic engineering for IP networks," *IEEE Network Magazine*, vol. 14, no. 2, pp. 11–19, Mar. 2000.

[5] Cisco, Enhanced Interior Gateway Routing Protocol [Online]. Available: http://docwiki.cisco.com/wiki/Enhanced_Interior_Gateway_Routing_Protocol.

[6] H. Newton, *Newton's Telecom Dictionary 24th Edition*. New York: Flatiron Publihing, 2008, p. 683.

[7] B. Fortz and M. Thorup, "Optimizing OSPF/IS–IS weights in a changing world," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 4, pp. 756–767, May 2002.