Macrothink Institute™

Network Protocols and Algorithms
ISSN 1943-3581
2010, Vol. 2, No. 4

Study and Performance of Interior Gateway IP Routing Protocols

Sandra Sendra, Pablo A. Fernández, Miguel A. Quilez and Jaime Lloret

Integrated Management Coastal Research Institute, Polytechnic University of Valencia.

C/ Paranimf nº 1, Grao de Gandía – Gandía, Valencia (Spain)

E-mail: sansenco@posgrado.upv.es, pabferm1@epsg.upv.es, mglngl14@gmail.com, jlloret@dcom.upv.es

Received: December 1, 2010 Accepted: December 31, 2010 DOI: 10.5296/npa.v2i4.547

Abstract

Large IP networks cannot be possible without routing protocols providing the appropriate paths between end sites. Many interior gateway routing protocols have been developed based on well known algorithms for IP networks in order to prevent routing loops. Some of these interior gateway routing protocols are RIP, OSPF and EIGRP. Each one of them has its benefits and drawbacks, and the best election depends on many parameters such as the network features, network hardware, few bandwidth wastage, scalability, costs, etc. In this paper we present a survey and a test performance of the main interior gateway IP routing protocols (although some of them can work with other non-IP protocols) that are used inside the Autonomous Systems. We will see which one provides the lowest delay path, the lowest number of hops, the lowest convergence time, and how the traffic sent through the network have different behavior depending on the routing protocol running in it, among others. Finally, we will conclude our paper providing some advises to the IP network designers and administrators. This work will help them to take the best election depending on their case.

Keywords: Routing Protocols, IGP, Autonomous System, RIP, OSPF, EIGRP, Performance Test.

1. Introduction

Working over large scale IP networks, moving between different sub-networks and reach the final destinations, is almost impracticable without using a routing protocol.

Since early 1980, the number and the size of the networks have increased hugely. Moreover, to find networks using different specifications and implementations is very common. This fact hampered information exchange between networks. In order to resolve this problem of incompatibility networks, the International Organization for Standardization (ISO) [1], developed a network reference model that would help manufacturers to build compatible networks with other networks. The reference model was called Open Systems Interconnection (OSI) [2]. It was a structured network model of several levels. Each one defined the architectures of communications systems interconnection and it included from the bit concept to the implementation concept.

This model is divided into seven layers and each of these layers specifies the functions and protocols to be used. These layers are (from up to down): Application, Presentation, Session, Transport, Network, Data Link, and Physical.

In this paper, we study the routing protocols, which are found in the network layer. The main aim of the network layer is to support the appropriate path from the origin to a destination, even though both are not connected directly. This layer is responsible for establishing, maintaining and terminating the connection. Also, this level defines the routing and controls the congestion packets in a sub-network. Devices that are assigned to this task are called routers. A router is a device that works at layer 3 (although in special occasions it could work in layer 2,) which is responsible for sending datagrams from one network to another network directly connected, and it is in charge of interpreting the routing protocols to carry information from origin to destination.

In a large network, we can distinguish different types of routing protocols [3]. On the one hand there is the Exterior Gateway Protocol (EGP). It is a protocol used to exchange routing information between autonomous systems. An EGP protocol allows knowing about all routes from its immediate neighbors on the internal network and the machines associated with them but lack of information about the rest of the Internet. Moreover, there is the Interior Gateway Protocol (IGP) that refers to the protocols used within a single autonomous system. An EGP determines whether a network is accessible from an autonomous system, while IGP will be responsible for dealing the routing within the autonomous system. Also, an autonomous system (AS) is a set of IP networks and IP devices that are managed by a single entity and have a common definition of Internet paths. To the outside world, an AS is a single entity that can be administered by one or more operators, while presenting a unified scheme for routing to the outside world [4]. Figure 1 shows an example of a network of three AS, where each AS uses an IGP protocol to communicate. While EGP protocols are used to communicate with other AS.

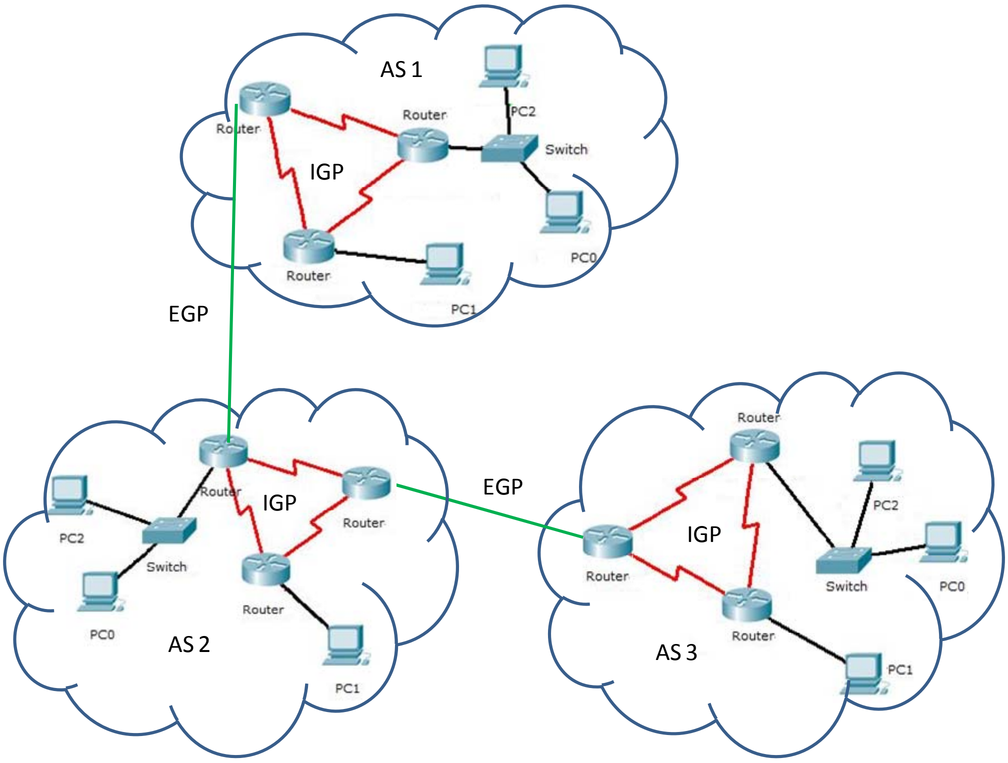

Figure1. Example of network of three AS

A routing protocol is a communication scheme between several routers that can form a sub-network. A routing protocol allows the router to share connected network information with other routers as well as its proximity to other routers. The information obtained from another router, by the routing protocol, is used to create and maintain the routing tables. The routing table contains the known networks and the interfaces associated with such networks.

The end user of a network does not know how data flows between the devices. The user desires to receive the information as soon as possible, no matter on what happens on the network. Moreover, the received information should be complete, without losses generated by network failures. These features are known as the fault tolerance of a network. The routing algorithm is responsible for determining a possible alternative route in the case of a link failure, or other topology changes, and for ensuring that the data structures are correct. This will also ensure the efficiency of their work. The routing protocol must maintain a balance between the design parameters such as scalability, fault tolerance, low overhead and QoS. But depending of the type of devices, other features should be added such as energy saving, resource limits, processing constraints, etc. A routing protocol must also show a high degree of accuracy, stability, robustness, fairness, simplicity and optimality. A good choice of these parameters allows working with real-time applications, and allows a very fast communication between devices.

This paper analyzes IGP routing protocols in detail, i.e., we will see the most well known routing protocols that can be used within a single AS. Initially, we will classify them depending on the mode of operation and on the way the routers update their routing tables.

Then, we will see the main features, advantages and disadvantages. They are obtained from their RFCs (Router Information Protocols (RIP) [5], Router Information Protocols V.2 (RIP V.2) [6], Open Shortest Path First (OSPF) [7] and Enhanced Interior Gateway Routing Protocol (EIGRP) [8, 9, 10]). The used metric, the type and the format of their packets and their operation will also be seen. Finally, we will show a set of performance tests over a network, in order to compare them quantitatively.

The rest of this paper is structured as follows. Section II shows the related work on routing protocols, separated by protocol type. Section III shows a brief classification of the protocols in terms of their mode of operation. An explanation of each protocol included in this study is shown in Section IV. A discussion about the main features of each protocol is shown in Section V. Section VI describes the software and hardware resources of the devices used to the test, the measurements performed and the graphs obtained for each routing protocol. Finally, Section VII concludes the paper.

2. Related Works

In general, there are very few works about routing protocols performance. In this section we have grouped some of these works by protocol type and, finally, we will show some studies that compare routing protocols.

One of the works that can be found in the literature about RIP is presented by C. Hedrick in [5]. This work is the RFC 1058 and provides a lot of information about this routing protocol. RIP is the oldest protocol. It uses a distance vector algorithm to form the routing tables and calculates the distance to a destination host in terms of how many hops a packet must traverse. It also shows technical aspects of the packet format and the metric. Due to its small number of hops, RIP is not created for large systems. Several methods have been added to the RIP protocol in order to solve some problems such as the generation of loops. Another work related to fault tolerance that uses RIP is [11]. In this work, D. Pei et al. show the design and development of a method for detecting RIP routing updates. Specifically, RIP-TP protocol is presented. It uses hop count as routing metric. The authors emphasize its efficiency, simplicity, low operating cost and compatibility with the standard RIP. In order to assess the design efficiency, they show a series of experimental simulations to demonstrate that it is possible the improvement of fault detection in routing protocols. They particularize these evidences with RIP.

There is published a work about OSPF routing protocol in [12]. It is presented by A. Shaikh et al. In this paper, its authors provided a study of the OSPF behavior in a large operational network, based on a hierarchical structure formed by 15 areas and 500 routers. One of its main features of this network is that it provides highly available and reliable connectivity from customer’s facilities to applications and databases residing in a data center. They introduced a methodology for OSPF traffic analysis, analyzing the link-state advertisement (LSA) traffic which is generated when the network experiments a topology change. Also the authors provide a general method to predict the rate of refresh LSAs from router configuration information and a set of measurements confirm that the method is accurate. Moreover, the authors observed that the type of topology could provoke certain asymmetries in duplicate-LSA traffic. Finally they showed a method for reducing duplicate-LSA traffic by altering the routers' logical OSPF configurations, without changing the physical topology of the network. Another study of OSPF is shown in [13]. This work, presented by A. Basu et al., studies the stability of the OSPF protocol under steady state and with interferences. In this study we will see what effects are given by the TE (Traffic Engineering) extensions on the stability of a network when OSPF is running. OSPF TE extensions provide mechanisms for ensuring that all network nodes have a consistent view of the traffic parameters associated with the network. The authors also analyze whether it is possible to accelerate the convergence time of the network, analyzing the Hello packets and the number of route flaps caused by a failure in the network, because the number of route flaps characterizes the intensity disruption of the network. The authors conclude the paper letting us know that the OSPF-TE protocol seems fairly stable, and adding that extensions TE does not significantly change the times of convergence, even in presence of multiple failures. But, a high number of failures in the network could lead to overload of the processor because it will have to attend a large number of alerts in the short term.

Because EIGRP is a Cisco proprietary protocol, sometimes, it is quite difficult to find information about it. B. Albrightson et al. let us know in [14] that EIGRP is based on IGRP protocol, but improving their benefits. They explain that EIGRP is a protocol based on a hybrid routing algorithm, sharing some properties of distance vector and link state algorithms. This protocol is the first Internet protocol that addressed the loop problem. Other aspects which shows are the type of metrics, the transport mechanisms and the methods used to discover the networks, among other features. A. Riesco et al. applied the EIGRP algorithm to an application based on Maude [15]. Maude is a programming language for formal specifications using algebraic terms. It is an interpreted language that allows the verification of properties and transformations on models that can run the model like a prototype. The authors show how to build an infrastructure of processes implemented by Maude, giving the chance to send a message directly to a neighbor or broadcast to all neighbors. EIGRP protocol implements the top of this basic infrastructure. Finally, the global system is tested and analyzed. The analysis is based on the search command that proves if a "bad" success could happen. This allows verifying the model, which examines whether a formula is true for all conditions.

We have also found some works and a master thesis that compare and make some kind of testing with multiple protocols.

There is a master thesis that shows a comparative analysis as the work (M. Nazrul and Md. A. Ullah in [16]). Their goal was to evaluate which protocol, EIGRP or OSPF, is most suitable to route in real-time traffic. The simulations are based on the convergence Time, Jitter, End-to-End delay, Throughput and Packet Loss. They demonstrated that EIGRP has faster convergence time than OSPF, because EIGRP can learn the topology information and updates faster than the others. Another important issue is that the packet delay variation for EIGRP is better than for OSPF, and consequently data packets in EIGRP reach faster to the destination compared to OSPF. Also, EIGRP, present less number of lost packets and a higher throughput than OSPF, when there is high link congestion.

Another work of the same authors where there is a comparative analysis of the routing protocols EIGRP and OSPF is shown in [17]. In order to evaluate OSPF and EIGRP's performance, their authors designed three network models configured with OSPF, EIGRP and a combination of EIGRP and OSPF and the three topologies where simulated using the Optimized Network Engineering Tool (OPNET) [18]. In this case, the protocols and the combined use of them are also analyzed in terms of convergence time, jitter, end-to-end delay, throughput and packet loss. The evaluation results show that, in general, the combined implementation of EIGRP and OSPF routing protocols in the network performs better than each one of them alone.

We have found a master thesis where their authors use a combination of different Routing protocols (E. S. Lemma at al. in [19]). They use OPNET to carry out the network simulations, using a combination of EIGRP&IS-IS, OSPF&IS-IS. The main aim of that paper was to configure multiple routing protocols on a selected network topology and analyze the performance improvement of the network. They based their comparison analysis on several parameters that determined the robustness of these protocols. In order to do it, their authors simulated five different scenarios on the same network in order to reveal the advantage of one over the others as well as the robustness of each protocol combination and how this can be measured. The selected protocols for each scenario were OSPF, EIGRP, IS-IS, OSPF/IS-IS and EIGRP/IS-IS. The results show that the use of combined protocols in a network, improve significantly the network performance.

The increased use of new technologies incremented the possibility of malicious attacks to our network, which could cause data loss, loss of privacy and even, eventually can lead to large monetary losses. Therefore, in [20], the authors examine the advantages and disadvantages of MD5 (Message-Digest Algorithm 5) authentication system compared to non-secure system when EIGRP, RIPv2, OSPF routing protocols are used. MD5 is a 128 bits cryptographic reduction algorithm that is widely used. The authors measure values of delay, jitter and network overhead, in both cases for all protocols, and conclude that the EIGRP protocol shows the lower overhead, even when the system is heavily overloaded.

Finally, another paper that looks interesting is the work presented by C-C Chiang et al. in [21]. As we know, the security mechanisms play an important role in networks and in the Internet world. There are many ways to find vulnerabilities in a network and launch attacks against the network. In this paper, the authors examine the performance and security problems of several existing routing protocols including RIP, OSPF and EIGRP. Several routing performance parameters are evaluated and analyzed through using SNMP (Simple Network Management Protocol) sessions. They briefly describe the three IGP protocols, their network Infrastructure and the experimental evaluation methods. In opposite of denegation of service (DOS) attacks and contaminated tables, which are among the most serious attacks to network topologies, the authors propose an automatic mechanism to analyze the states of routing and intrusion detection in real-time response. The study concludes that the distance vector routing protocols are more robust than link-state routing protocols for the unstable network topology because global link-state's flooding of updates increase when link state changes. But, the distance vector algorithms can only used for small networks.

3. IP Routing Protocols Classification

As we have defined, a routing protocol is an algorithm procedure to obtain a set of instructions or routes to take from one network to another. These instructions may be provided to a router from other dynamically, or can be statically assigned to the router by an administrator. The network administrator takes into account many aspects when he/she selects a routing protocol:

- The size of the network
- The bandwidth of the available links
- The processing capacity of the routers
- The manufacturer and models of the routers
- The protocols that are already in use on the network

So, the protocols can be classified as follows [22]:

Non-adaptive algorithms: They are also known as static routing. Devices operating under these conditions do not base their routing decisions on measurements or traffic estimations. The route is previously calculated off-line and loaded into the routers when the network is set up. When using static routing, the network administrator has to configure manually the information about remote networks in each router. The main problem in this case is that changes are not possible in case of situations where there is traffic congestion or a failure. Furthermore, they do not have the scalability or capacity to accommodate the growth provided by dynamic routing. For this reason, if there is a change in the topology, the administrator must add or delete the static routes affected by such changes. On the contrary, in small networks with few changes, static routes require very little maintenance. In addition, static routes are used as backup routes (only will be used if dynamic routes fail) and it is the preferred method for maintaining routing tables when there is only one route to a destination network.

Adaptive Algorithms: They are also known as dynamic routing. They change their routing decisions to reflect the topology and adapt to the traffic changes. In a large network, the manual maintenance of routing tables may need an enormous management time. In complex networks, adaptive routing will avoid congestion in some parts of the network and improve its performance. When routers use dynamic routing, they carry out the routing operations using the information received from other routers about the network topology to evaluate the routes to each network. Such algorithms must be simple because will be processed in real time inside the routers taking care of their CPU and memory resources. They differ from the static algorithms on how they obtain the information of the topology and on the metric used. Dynamic routing protocols learn all the available routes. They include the best routes in the routing tables and discard the paths that are no valid. When there is a change in the network topology, the information known about the network must also change. This information should reflect an accurate and consistent view of the new topology. When all routers on a network know the same information, the network has converged.

Most dynamic IP routing algorithms fall into one of the two categories: Distance vector and link state), although a hybrid version mixing both exist [3].

Distance vector routing protocol: This protocol calculates the direction (vector) and distance to any network. It also is known as Bellman-Ford algorithm. The algorithm collects information about the distances for each network. This is measured in number of hops or in the number of router located in the path to a destination network. The distance vector algorithms do not allow a router to know the exact topology of the network. The router only uses hop count to determine the best route. A maximum number of hops value will be used to consider an unreachable network. The routing tables include information about the cost of each route (defined by its metric) and the IP address of the first router on the route to each network listed in the table. There are regular updates of the routing tables between routers.

Link State routing protocol: The link state protocol use Dijkstra algorithm (also known as Shortest Path First algorithm). It maintains a complex data base, also called as link state database, which contains full information about the remote routers and the exact network topology. A router, configured with a link state algorithm, generates information about itself, its connected links and the link status. This information is transmitted from one router to another and each router makes a copy of it, but never changes it. The data is transmitted using the Link State Advertisements (LSAs), also called Link State Packets (LSPs) to each neighbor router. The goal of these algorithms is to provide identical information about network connections to each router, so each router can calculate the best routes to each network.

4. Routing Protocols under study

This section describes the most used Interior Gateway Protocols.

4.1 Router Information Protocols (RIP)

4.1.1 Introduction to the protocol

RIP is the oldest routing protocol. It is based on distance vector algorithm (Bellman-Ford). In its first steps was used by ARPANET in 1969. Gradually, the protocol was gaining popularity and different implementations arose, so standardization was necessary. This standardization was carried out by Charles Hedrik in 1988. It can be found formally described in RFC 1058 [5]. Over the years RIP has evolve to version 2. RIP version 2 (RIPv2) defines certain improvements in functions and extensions of RIP version 1 (RIPv1). Probably the most important difference between them is that RIPv1 is a classful protocol while RIPv2 is a classless protocol. It allows RIPv2 to be more compatible with the modern routing needs. RIPv2 is formally described in RFC 1723, which was updated in the STD 56 (RFC 2453 [6]).

4.1.2 Metrics

RIP uses the hop count as a metric to select the best route to each network destination, i.e., it simply counts how many routers (hops) a message must traverse to reach its destination. Expression 1 shows this relationship:

$$M = d \tag{1}$$

Where d is the minimum number of hops between the source and the destination. The maximum number of hops allowed is 15, all routes with a greater distance are considered unreachable. Obviously, this is not the ideal way to handle situations in real time, because it does not take into account neither the delay, nor reliability, or the bandwidth. Even so, thanks to its simplicity, this protocol is compatible with the vast majority of manufacturers. It is an ideal solution for networks from small to medium size.

4.1.3 Message types

In order to maintain an updated routing database, routing protocols based on distance vector algorithms usually send update messages periodically. RIP sends update messages each 30 seconds. This is one of the drawbacks of RIP, because it has to update constantly the routing database and it generates a large amount of traffic on the network.

RIP update messages are encapsulated in UDP (User Datagram Protocol) segments. The UDP source and destination port is 520. In this case, the IP header and the data link headers use broadcast destination addresses. Fig. 2 shows the format of the RIPv1 message.

0	8	16 31
Command	**Version**	**Must be zero**
Address family identifier		**Must be zero**
IP address (network address)		
Must be zero		
Must be zero		
Metric (hops)		

Figure 2. RIPv1 message.

The header consists of the following fields:

- **Command.** It specifies the message type (1 for an application, 2 for a response)
- **Version.** It allows to differentiate RIP versions (1 for RIPv1, 2 for RIPv2)

It can include up to a maximum of 25 route entries. Each route entry consists of:

- **Address family identifier.** 2 for IP addresses and 0 to request the full routing table.
- **IP address.** It determines the direction of the target path.
- **Metrics.** It keeps track of the hops (between 1 and 16).

At first, the fields named "must be zero" were added to support the increase of spaces for IP addresses in the future, but most of these spaces have been used in RIPv2.

The main limitation of RIPv1 is that it is a classful routing protocol, so it cannot use VLSM (Variable Length Subnet Masking) or CIDR (Classless Inter-Domain Routing). RIPv2, which is a classless routing protocol, is born to bridge this gap. While RIPv2 maintains the same basic message format of RIPv1, it adds two important extensions. Fig. 3 shows these innovations.

0	8	16	31
Command	**Version**	**Must be zero**	
Address family identifier		**Routing label**	
IP address (network address)			
Subnet mask			
Next hop			
Metrics (hops)			

Figure 3. RIPv2 message.

The first innovation is the Subnet mask field because it allows adding a 32-bit mask in the RIP route entry. The second added innovation is the next hop field. The next hop address is used to identify a next hop router address better than the address of the sender router, if it exists. If the field is set to zero (0.0.0.0), the address of the sender router is the best next hop address.

Another improvement of RIPv2 is that the updates are sent using the multicast address 224.0.0.9. This feature lets RIP to consume less network bandwidth.

In addition, multicast updates require less processing than the not RIP-enabled devices. With RIPv2, any device that is not configured for RIP, discard the frame at the data-link layer, with the consequent processing savings. With this simple improvement, the problem of traffic generated is attenuated by the continuous update messages.

4.1.4 Tables

Routers running routing protocols that are based on distance vector algorithms must keep constantly updated their routing databases in order to perform the right decisions to find the best route. These databases include information about the total cost of each path and the next hop to reach each network. If the routers are not aware of the changes in the network, they may route packets to interfaces that are no longer connected to the best route.

4.1.5 Algorithm operation

When a router receives an update message with a change in it, it compares the distance of the new message for that destination network with the distance in the current routing table, if the new distance is lower than the existing one, the new route is accepted and the database is updated in order to use the new route. Then, the metric value is increased by 1 and the interface where the update came from is used as the next hop in the routing database. RIP routers only maintain the best route to a destination, but they can keep more than one route to the same destination if there is the same cost for those routes.

One of the main problems when this algorithm is running is that it need too much time to converge, and, as a result, there could be inconsistent routing entries, which it can cause routing loops. In order to prevent incorrect routing information, RIP implements split-horizon and standby timer mechanisms. The split-horizon rule is a mechanism that simply prevents sending information of a route through the same interface through which it learned that information. Standby timers are used to define the amount of time that a failed route must not be published. When a route fails, the standby timer is activated. During this period, a route with a metric greater than the original is not accepted. If the original route becomes active, or announces a route with a metric less than the original, it is accepted immediately.

These mechanisms reduce routing loops, but they do not eliminate them. Despite of all mechanisms, in order to avoid routing loops, RIP sets the hop limit to 15. If a packet reaches this value, the distance is considered infinite and it is discarded.

4.2 Open Shortest Path First (OSPF)

4.2.1 Introduction to the protocol

The Open Shortest Path First (OSPF) routing protocol is a public (open standard) that is based on the link state. RFC 1583 contains a description of the concepts and operations of OSPF Link State. A second version has been published n RFC 2328. The routing protocol calculates optimal routes based on lower cost of the links to a destination using Dijkstra's algorithm.

4.2.2 Metrics

Each router measure the cost to each one of its neighbors. The cost is used as a metric to determine the best route. A cost is associated with each router interface. Generally, the cost of a route is calculated using the formula shown in expression 2:

$$Cost = \frac{10^8}{\text{bandwidth (bps)}} \tag{2}$$

The lower cost, the more likely that interface will be used to send data traffic. Table 1 shows different cost values as a function of the type of link.

Table 1. Cost of several type of links.

Type of link and Bandwidth	Cost
Link serial de 56Kbps	1785
Link serial T1 1.544Mbps	64
Link serial E1 2.048 Mbps	48
4 Mbps Token Ring	25
Ethernet de 10 Mbps	10
16 Mbps Token Ring	6
100 Mbps Fast Ethernet, FDDI	1

4.2.3 Message Types

Each router sends hello packets to its neighboring routers in order to track their status. OSPF routers must have the same hello and dead time intervals to exchange their information. By default, some manufacturers, such as Cisco and Juniper, set the dead time interval three or four times the hello time interval (depending on the type of network, broadcast or non broadcast). In the case of four times, it means that a router has four or three chances to send a hello packet before being declared dead. The exchange of LSA is triggered by an event in the network instead of regular updates (this speeds up the convergence process). The Hello packets are sent every 10 seconds by default in multi-access broadcast networks and point to point. Generally, for interfaces that connect to NBMA networks, such as Frame Relay, the default time is 30 seconds to send the hello packets.

The messages used by OSPF are shown in Fig.4 and Fig. 5. Fig.4 is the header used in a hello packet, while Fig.5 is the hello packet with its data fields.

0	8	16	31
Version	**Type**	**Packet length**	
Router ID			
Area ID			
Checksum		**Authentication Type**	
Authentication data			

Figure 4. OSPF hello packet header.

0	8	16	31
20 Bytes OSPF HEADER			
Netmask			
Hello Interval		**Options**	**Router Priority**
Dead interval			
Designated Router			
Backup designated router			
Neighbor router ID			
Neighbor router ID			
You can add additional fields of the neighbor router ID of the end of the header			

Figure 5. OSPF hello packet.

OSPF topology structures the network into areas. Areas with more than 50 routers are not recommended. OSPF use LSP (Link State Packet, also known as LSA Link State Advertisement) to build and update its list of neighbors with the cost to them. After that, it floods LSP to all routers of the network in order to build the complete topology.

LSPs are only sent when there are changes. These packets are listed sequentially for each shipment. LSP have limited lifetime, in order to avoid having lost packets traveling through the network indefinitely. OSPF does not send a LSP through the interface where it had been received. It does not send a LSP if it has been sent before to avoid duplicity.

Each router updates its link state and topological database using LSPs. Once they have assembled all the information, each router calculates the best routes to all destinations in the network using SPF algorithm. It builds the tree topology with the router as root, and all the possible routes to each network. The link-state routing algorithm is fully aware of distant routers and how they interconnect. After running SPF, the routing table is built.

OSPF use seven different types of messages. These messages are used to establish the adjacencies with the neighbors and to exchange information. Due to the different types of LSA messages, OSPF is so scalable.

Moreover, it is possible to define different types of areas using LSAs or other messages. The standard area is responsible for accepting Link State Updates (LSUs) and route summarization. Backbone area (area 0), is a transit area. All areas must be connected to it. It has the same features than the standard area. Stub area and totally stubby area do not accept external routes (from another autonomous system or another area). In order to reach these areas, a default route has to be defined and it has to be redistributed throughout the area. Not-so-stubby area (NSSA) is similar to the stub area but allows you to import external routes.

4.2.4 Node status

The relationship between neighbors can be in seven states:

- **Down state:** Routers do not exchange any information. They are waiting to start the Init state.
- **Init state:** Routers send hello packets (usually every 10 seconds). When a hello packet is received from another router, there could be a neighbor.
- **Two-Way state:** Hello packets include the router ID and the ID of the discovered neighbors. Two-Way state is started when the router receives a Hello packet which includes its ID.
- **ExStart state:** This is a transient state to the next state (Exchange State) in which Hello packets are exchanged to determine who will be the Master and who will be the Slave.
- **Exchange state:** Now the routers exchange Database Description packets (BDPs) which include a summary of the link state database (topological database).
- **Loading state:** This state is where the routers make requests for the new information. For this purpose Link State Request (LSR) packets are used.
- **Full Adjacency:** When the Loading state finishes, with the arrival of all LSAck, the router begins the state of Full Adjacency.

Routing updates produce a large volume of traffic when there are topology changes. Every time a LSA packet causes a change in the link-state database, SPF algorithm recalculates the best paths and updates the routing table. For this reason, to limit the amount of link-state routers within an area is necessary.

OSPF supports VLSM, Complex metrics and multiple routes (although the routes selected may not be symmetrical).

4.2.5 Algorithm operation.

Many routers can be connected to a broadcast multi-access network segment. If every router has established the full adjacency with all other routers and exchange link state information with them, they could need too much process capacity. In this case, expression 3 gives the number of adjacencies needed when there are n routers.

$$\#Adjacencies = \frac{n\,(n-1)}{2} \tag{3}$$

In order to reduce the number of adjacencies to exchange routing information, OSPF routers select a designated router (DR) and a backup designated router (BDR) that serve as central points to update the routers and exchange routing information. When OSPF priorities are the same, the DR election is taken by the router ID. The highest router ID is selected. The router with the second highest priority will be BDR. In peer to peer networks there are only two nodes, so there is not DR election. Both routers become fully adjacent to each other.

DR sends link state information to all other OSPF routers in the same segment through the 224.0.0.5 multicast address. Although the efficiency is higher than when there is not a DR, the disadvantage is that now there is a single point of failure. In order to ensure that both the DR and the BDR see the state of the other routers they use the 224.0.0.6 multicast address to receive the updates from the other routers.

4.3 Enhanced Interior Gateway Routing Protocol (EIGRP)

4.3.1 Introduction to the protocol

EIGRP is a hybrid routing protocol, owned by Cisco Systems. Cisco EIGRP was launched in 1994 as a scalable and improved version of IGRP (which is also a proprietary distance vector routing protocol). Therefore there is no associated RFC.

EIGRP improves the convergence properties and operates more efficiently than IGRP. It makes efficient use of bandwidth. It uses a minimum bandwidth when the network is stable. EIGRP routers do not send the full routing tables, but send partial and incremental updates. The maximum number of hops in EIGRP is 224. It also supports CIDR and VLSM, which allows network designers to maximize the address space.

EIGRP uses RTP (Reliable Transport Protocol) as transport layer protocol to guarantee the delivery of routing information. With RTP, EIGRP can send multicast and unicast packets simultaneously with higher efficiency.

EIRGP is able to route IP, IPX and AppleTalk.

4.3.2 Metrics

EIGRP uses 32 bits long metric. This metric can be expressed as shown in expression 4.

$$M = \begin{cases} \left[K_1 * BW + \frac{K_2 * BW}{256 - load} + (K_3 * delay)\right] * \left[\frac{K_5}{reliabylity + K_4}\right], K_4 \neq 0 \; and \; K_5 \neq 0 \\ K_1 * BW + \frac{K_2 * BW}{256 - load} + (K_3 * delay), \qquad K_4 = K_5 = 0 \end{cases} \tag{4}$$

The default values are shown in table 2.

Table 2. Constant Values.

Constant name	Value
K_1	1
K_2	0
K_3	1
K_4	0
K_5	0

Using the default values shown in table 2, the metric can be expressed by expression 5.

$$Metric = BW + delay \tag{5}$$

4.3.3 Message Types

IGRP relies on different type of packets to keep its tables and establish relationships with neighboring routers. The five EIGRP packets are:

- Hello Packet
- Acknowledgement Packet
- Update Packet
- Consultation Packet
- Reply Packet

Fig.6 shows the EIGRP Packet Format [23].

<table>
<tr><td>0</td><td>4</td><td>8 … 31</td></tr>
<tr><td><b>Version</b></td><td><b>OPCode</b></td><td><b>Checksum</b></td></tr>
<tr><td colspan="3"><b>Flags</b></td></tr>
<tr><td colspan="3"><b>Secuence</b></td></tr>
<tr><td colspan="3"><b>Acknowledgment</b></td></tr>
<tr><td colspan="3"><b>Autonomous System number</b></td></tr>
<tr><td colspan="3"><b>TLVs</b></td></tr>
</table>

Figure 6. EIGRP Packet.

Where each field of bits has the following meaning:

- **Version:** Specifies the version of EIGRP. Version 2 is the most recent version.
- **OPCode:** Specifies the type of EIGRP packet. Opcode 1 is the update packet, opcode 3 is the Query, opcode 4 is the reply, and opcode 5 is the EIGRP hello packet.

- **Checksum:** It is used as the regular IP checksum. It is calculated based on the entire EIGRP packet, excluding the IP header.
- **Flags:** Involves only two flags. The flag indicates either an init for new neighbor relationship or the conditional receive for EIGRP RTP.
- **Sequence:** Specifies the sequence number used by the EIGRP RTP.
- **Acknowledgment:** Used to acknowledge the receipt of an EIGRP reliable packet.
- **Autonomous System Number:** Specifies the number for the identification of EIGRP network range.

Routers establish relationships using Hello packets. These packets are sent by the router every 5 seconds. When a router receives a packet Hello a new neighbor is found. When EIGRP routers form adjacencies with other routers, it can dynamically learn new routes from the network. They also identify networks that became unreachable or inoperable, and can detect routers that had not been reached for a period of time.

The hello packet contains the hello and dead time interval. The dead time interval is three times the hello time interval.

4.3.4 Tables

EIGRP uses three tables for its operation: neighbor table, topology table and routing table. The neighbor table is the most important. Each EIGRP router contains a table of its neighbors. Each table entry contains the address of the interface of the router, the timeout, the round trip time timer (SRTT), the queue number (Q Cnt) and the sequence number (Seq No). This information is stored in the neighbor data structure.

4.3.5 The algorithm operation.

The core of EIGRP is the finite state machine (FSM) Diffuse Update Algorithm (DUAL) which is the route estimation engine. The DUAL FSM contains all the logic used to calculate and compare routes in an EIGRP network. DUAL tracks all routes published by the neighbors, then it compares the routes through the composite metric for each route and ensures that the route is free. It adds the lowest cost routes in the routing table and the places of the successor route. DUAL uses the information of these tables, to quickly select alternative routes. If a link is disabled, DUAL seeks for an alternate route in the topology table, ensuring loop-free operation, and allowing simultaneous synchronization between all involved routes.

The topology table contains all the EIGRP routing tables and the EIGRP tracks this information so EIGRP routers can identify and switch to alternative routes. The information received from DUAL is used to determine the path of a successor, which is the term used to identify the best option. This information is added to the topology table. A successor is a route selected as the main route to reach a destination and it is possible to have up to four successor routes for each destination. These can be equal or unequal cost and can be identified as the best routes without loops.

A feasible successor (FS) is a backup route. These routes are identified at the same time as successors, but only remain in the topology table. There could be multiple feasible successors for a destination in the topology table. If a path breaks down, the router looks for a feasible successor. If it is not possible to identify a feasible successor based on current information, the router puts an active state and sends route query packets to all neighbors to recalculate the current topology.

If a hello packet is not received within the dead time interval, the time expires and the DUAL estimates the new topology.

5. Routing protocols comparison.

Over time, routing protocols have evolved and have been adapted to the needs of increasingly complex networks. Therefore, when a protocol is chosen, many factors must take into account. Some of them are the type of network it is wanted to be implemented, which future applications will be running in the network, its potential growth, etc. For example, if a simple routing protocol is needed to set up a small network and it must be supported by most routers of the market, we should use RIPv1. However, RIPv1 propagates updates using broadcast packets, so there could be significant bandwidth consumption and RIPv1 does not support any authentication mechanism, so there could be security holes.

In order to avoid some limitations of RIPv1, for example classful addressing, a network administrator can use RIPv2, which allows VLSM and CIDR. Moreover, RIPv2 supports authentication. This characteristic ensures that routers only accept routing information from routers that have been authenticated.

Moreover, if there is a large network and we need fast convergence times, we will undoubtedly use OSPF or EIGRP (we will see these measurements later). However, the design of these routing protocols in an IP network is more complex, and the configuration complexity and the administration time is higher.

When the network hardware is from different manufacturers, the administrator must use open routing protocols, so in this case EIGRP cannot be used, only RIPv1, RIPv2 and OSPF could be used.

Table 3 shows the main features of the routing protocols RIPv1, RIPv2, OSPF and EIGRP.

Table 3. Main features.

Feature	RIPv1	RIPv2	EIGRP	OSPF
Algorithm	Bellman-ford	Bellman-ford	Diffuse Update Algorithm (DUAL)	Dijkstra
When it does updates?	30 Sec.	30 Sec.	when there is a link failure	when there is a link failure

What is the Metric based on?	Hops	Hops	Bandwidth, Delay, load and reliability	Bandwidth
A designated router is selected in a multi-access network?	NO	NO	NO	YES
Updates using Multicast addresses?	NO	YES	YES	YES
Support VLSM	NO	YES	YES	YES
Protocol type	Distance Vector	Distance Vector	Hybrid Routing (Distance Vector & Link state)	Link state
Authentication Support	NO	YES	YES	YES
Route summarization Support	NO	YES	YES	YES

6. Performance Test

In order to test the efficiency of the protocols previously detailed, we have set up topology formed by different models of the Cisco Systems manufacturer. Several tests have been made and the results are shown in several graphs. These results will demonstrate what protocol has the best behavior.

In this section, first we will describe the devices used and the type of media used to set up the topology. We will also explain the program software used. Finally, we will describe each test performed and the results obtained.

6.1 Hardware and software resources

This subsection describes the software and hardware resources used to perform our test bench. Furthermore, the topology and the IP addressing configuration used in these tests is shown.

In order to measure the performance of each protocol, we used a network topology of 12 routers forming a random topology which has multiple paths to go from one side of the topology to the other. The routers used in our test bench are the following models: Cisco 1720, Cisco 1721 and Cisco 2611 router. All these routers use a Motorola PowerQUICC MPC860 processor to run the operating system and execute the coded instructions and its subsystems. It has enough processing capacity to perform the basic operations needed to accomplish the functionality of the router. All of them have two slots for serial WAN Interface Cards (WIC), which support synchronous and asynchronous serial links, and can support different protocols such as Point-to-Point Protocol (PPP) or Frame Relay and Data

Terminal Equipment/Data Communications Equipment (DTE/DCE). DRAM is used to run the Cisco IOS software and its subsystems, the routing tables, the fast switching cache, the running configuration, and so on. It has several slots for installing additional RAM memory. The main features of these devices are shown in table 4.

Table 4. Devices used.

	Devices		
	Cisco 1720 Router	***Cisco 1721 Router***	***Cisco 2611 Router***
Processor	Motorola MPC860T PowerQUICC at 48 MHz	Motorola MPC860P PowerQUICC at 80 MHz	Motorola MPC860 40 MHz
DRAM Memory	Onboard: 16 MB	Onboard: 64 MB	Onboard: 64 MB
Flash Memory	4 MB	32 MB	8 MB

In addition, a Switch Cisco Catalyst 2950 with 24 Fast Ethernet ports has been used in the centre of the topology. We have also used 3 Personal Computers as the end devices. PC_A and PC_B are used to send traffic through the network. PC_C is used to measure the traffic that travels through the network.

Figure 7 shows the network topology used in the test bench and the IP address of each router interface.

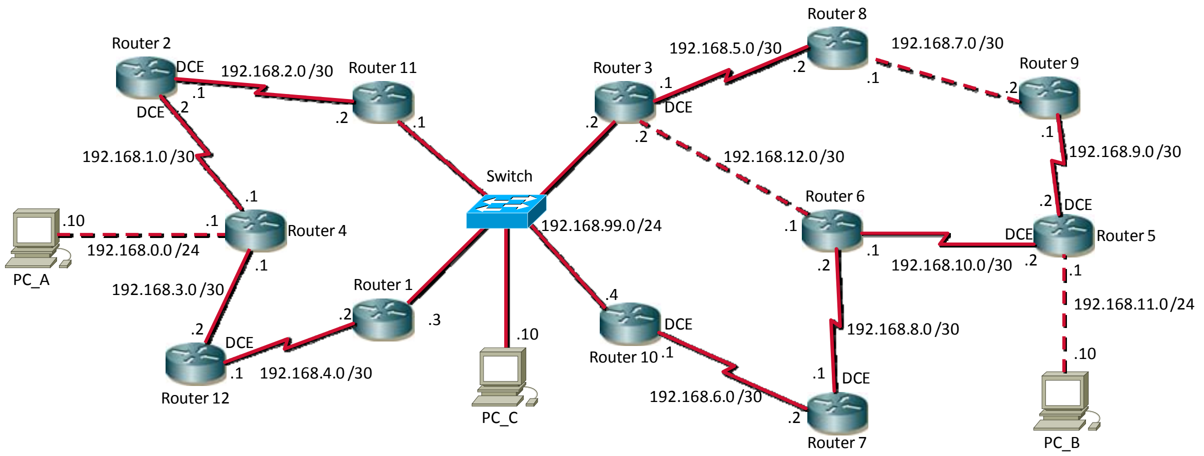

Figure 7. Network topology

Fig. 8 shows the legend of the type of connections used in the network topology

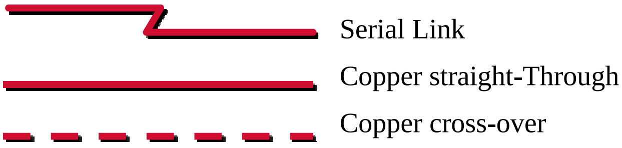

Figure 8. Cable Legend

PC_C is connected to the central Switch. This device is configured in monitor mode [24], in order to monitor the traffic that crosses it. Furthermore, we have limited the bandwidth in the serial links to 128 Kbps to have significant differences between paths.

In order to take our measurements we have used the following commands and software applications:

- **Ping command**. It is a command that is available in almost all operative systems. It helps to determine end to end reachability at IP and it is used to verify if a network data packet is capable of being distributed to an IP address without errors. So it allows us to see if the remote computer is reachable and is able to send information back. Moreover, it provides the round trip time, which will be used by us to know the IP routing protocol differences.

- **Net Meter** monitors network traffic and shows the bandwidth rate used by a network device. Net Meter is very small software that does not install extra drivers to the computer. It monitors network traffic through all network connections on the computer and displays real-time graphical and numerical downloading and uploading speeds. The program can displays transfer rates of multiple network connections at the same time and it can also record transfer rates of connections in text or Microsoft Excel CSV format [25].

- **Network Associates Sniffer PRO** is a tool that makes statistical analysis of the network activity. It is not a sniffer that provides the packets that traverse the network. This software is designed to help us to see the network traffic, to control the network and displays all this information as a function of the time. The measurements that can be shown are the use of the communication channel, packets/s, bytes/s, lost packets, broadcast and multicast packets, errors and collisions in the transmission, among other measurable parameters [26]. It provides the following capabilities:
  - Capture and decode data on a network.
  - Analyze the activity involving specific protocols.
  - Generate and displays statistics on network activity.
  - Perform analysis of patterns of the network activity.

6.2 Convergence time comparison

In order to perform this test, we followed next procedure. First, we wait until the network is stable and converged. Then, we check the route between PC_A and PC_B in order to know which path was followed for each routing protocol. For all protocols the path between PC_A and PC_B was router4-router12-router1-router3-router6-router5. Then, we sent a continuous ping between PC_A and PC_B in order to measure how many time was no transmission through the switch when the router1 is down. Next, we waited until the network converged for all routing protocols and gathered each convergence time. We saw that the path for this time was router4-router2-router11-router3-router8-router9-router5. Finally, we measured a second failure in the network. This second time was caused in order to know how their performance was when recovering a second time. In this case, the failed router was router3. We also waited until a path between PC_A and PC_B was recovered and the information arrived again to the destination. The final path was router4-router2-router11-router10-router7-router6-router5.

Fig.9 compares the convergence time of each protocol under our study. The time is expressed in milliseconds. The "1" value represents the convergence of the network, so the connectivity between the two end points exist. While "0" represents that the network is not converged yet.

As we can see, EIGRP is the protocol that first converges (in approximately 5 seconds). It is followed by OSPF. The protocol that needs more time is RIPv1 (it needs around 27 seconds).

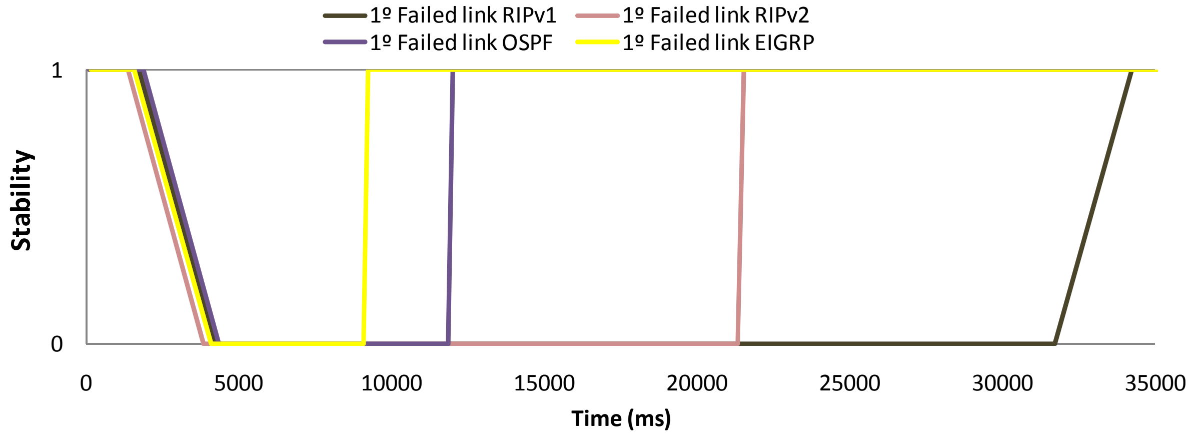

Figure 9. Convergence time comparison

Once we have already tested and compared the convergence capability of each protocol, we cause a second consecutive failure in the network in order to check their robustness to successive failures and their process capacity to find another alternate route. The second test will show the time taken by each protocol when the second link fails down. It is shown in Fig.10. The time is expressed in milliseconds. "1" value indicates that the network has converged, while "0" indicates that the network has not converged.

As we can see, again, the fastest protocol is EIGRP. OSPF and RIPv2 have similar convergence times. In comparison of RIPv1, EIGRP is approximately 3 times more rapid than RIPv1.

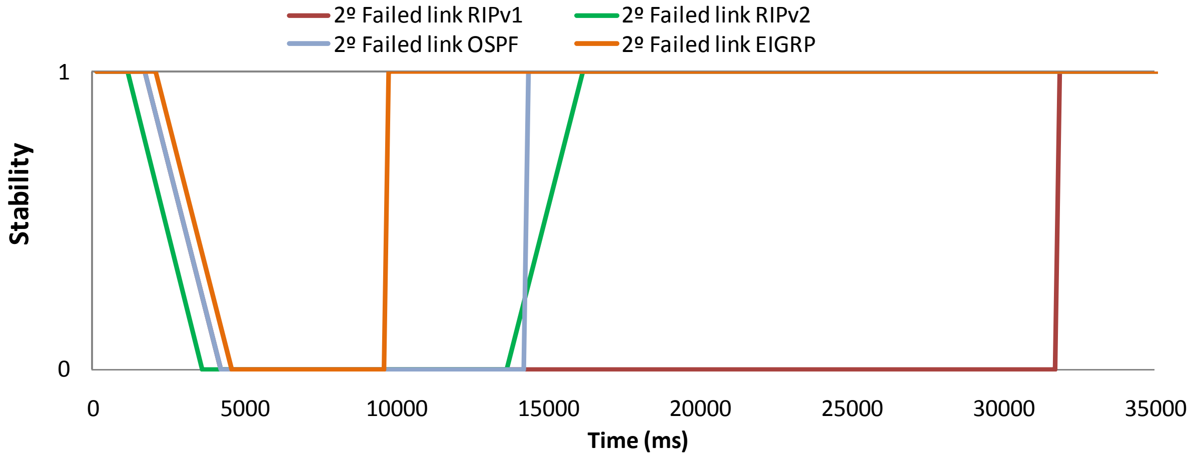

Figure 10. Convergence time comparison when a second link fails down.

6.3 Bandwidth consumption comparison

In order to test the maximum bandwidth consumption through a path, we sent a large file (about 1.36 GB) between PC_A and PC_B and vice versa, in order to assure that the behavior was the same in both cases, while running each routing protocol. The maximum bandwidth consumption is fixed by the maximum speed provided by the link with lowest maximum bandwidth. In this case it has been limited by the serial links bandwidth (which is 128kbps).

In this case we used the Netmeter in PC_A and PC_B to take the measurements.

Fig.11 shows the downlink bandwidth of each routing protocol at PC_A. As we can see, the behavior of all protocols is very similar and they range between 0 and 41.200 bps, with the exception of OSPF, whose bandwidth consumption is not regular, there are 15 seconds of low bandwidth consumption every 45 seconds. We can observe that the file is completely sent. Their average bandwidths have been EIGRP=23790 bps, OSPF=21037 bps, RIPv1=25.052 bps, RIPv2=23.340 bps. This means that the worst routing protocol was RIPv1 because is the one that consumes more average bandwidth and the best is OSPF.

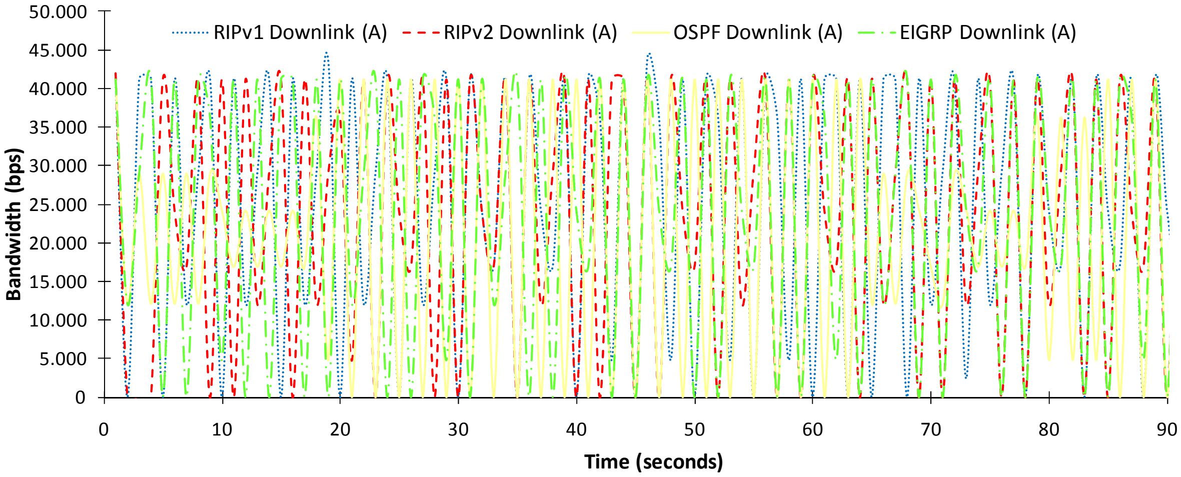

Figure 11. Downlink bandwidth in PC_A

Fig.12 compares the bandwidth consumed at the uplink of PC_A when each routing protocol is running. Fig. 12 shows that EIGRP and RIPv2 have the most stable behavior. With these protocols there is an average bandwidth consumption around 42.000 bps. On the other hand, RIPv1 presents a down peak every 30 seconds approximately, which is the time when it generates the updates. Finally OSPF shows the lowest average bandwidth consumption, due to their constant fluctuation between 12000 and 41200 bps (with its lowest peak in 113 bps) that locates its average value in 33911 bps.

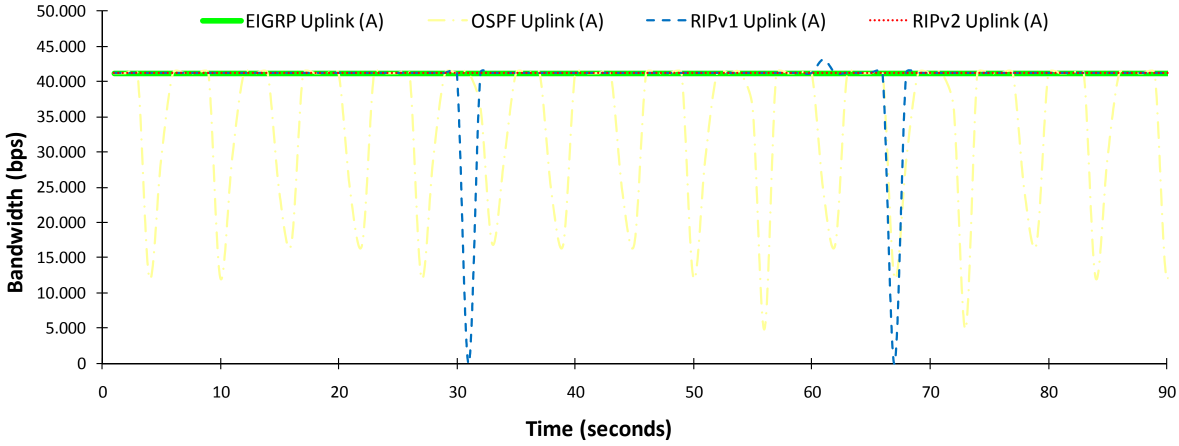

Figure 12. Uplink bandwidth in PC A

Fig.13 compares the bandwidth consumption in the downlink case for all routing protocols. Their behavior is quite similar than the one obtained in Fig.11. RIP2 and EIGRP present the most stable bandwidth consumption, which is approximately 41200 bps. RIPv1 and OSPF have periodical fluctuations that generate lower average bandwidth consumption. In the case of OSPF the average value is 34316 bps.

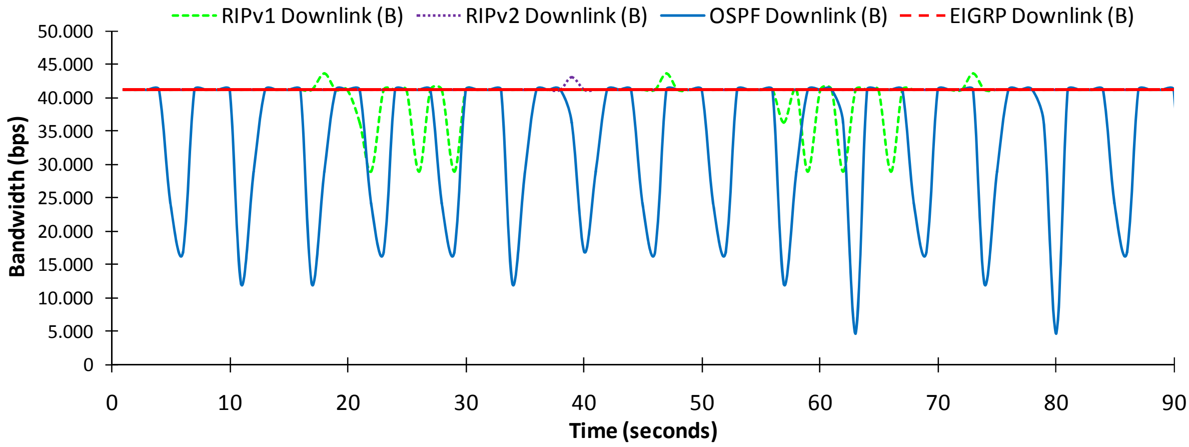

Figure 13. Downlink bandwidth in PC_B

Fig. 14 shows the comparative of the uplink bandwidth consumption for each routing protocol in the other endpoint of network (PC_B). In this graph there is not any strange behavior, all the protocols have a fluctuation between 0 and 41200. The average value for each routing protocol has been RIPv1=25178 bps, RIPv2=23804 bps, OSPF=21058 bps, EIGRP=23804 bps, so the lowest bandwidth consumption has been OPSF.

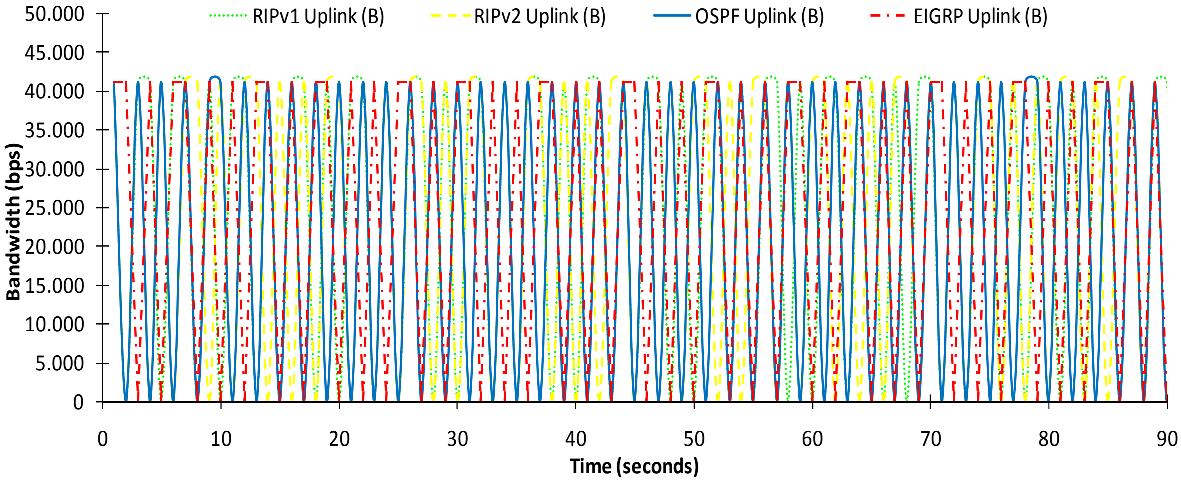

Figure 14. Uplink bandwidth in PC_B

6.4 Comparison of the # of packets sent to the network

In order to study the behavior of the network for each routing protocol, we have measured the number of packets through the network when a large file is sent between both endpoints (PC_A and PC_B). This information has been measured by Sniffer Pro in PC_C. The file had a size of 1.36 GB.

Fig. 15 shows the comparison of the packets sent from PCA to PCB when each routing protocol is running. There, we can see the number of packets along the time during 280 seconds. As we can see, RIPv1 and EIGRP are the most unstable protocols, showing high differences between maximum peaks and minimum peaks, while RIPv2 and OSPF prove to be more stable. RIPv1 had the lowest average value (6.79 packets per second), while EIGRP had the highest average value (9.64 packets per second). EIGRP graph had many peaks with more than 15 packets per second. RIPv2 and OSPF had 8.86 and 8.50 packets per second respectively. RIPv1 and RIPv2 are the slowest ones in stabilize. They took around 50 or 60 seconds. EIGRP is the fastest to be stabilized (it took 10 seconds).

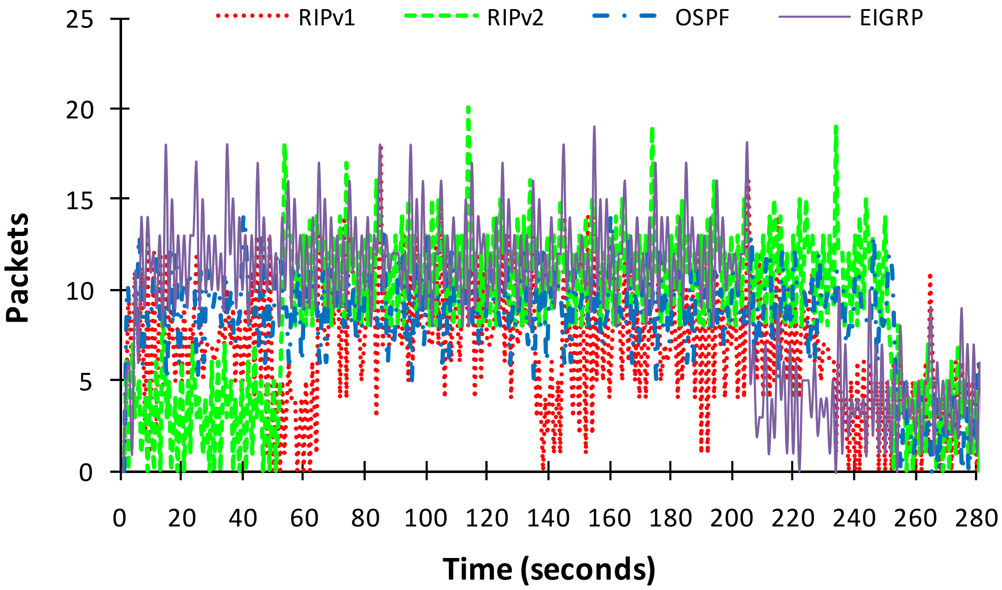

Figure 15. # of packets from PC_A to PC_B

Fig. 16 shows the comparative of the packets sent from PC_B to PC_A when each routing protocol is running. In this case all protocols have a similar behavior. All of them reached their stability between time=20 seconds and time=40 seconds. Again, EIGRP protocol has been the one with highest average number of packets per second (7.46 packets per second). This time RIPv1 had 7.27 packets per second, RIPv2 had 6.53 packets per second and OSPF has been the lowest one with 5.15 packets per second. RIPv2 and OSPF protocols are quite stabilized until the 220 seconds, while RIPv1 and EIGRP end their stability at around 240 seconds.

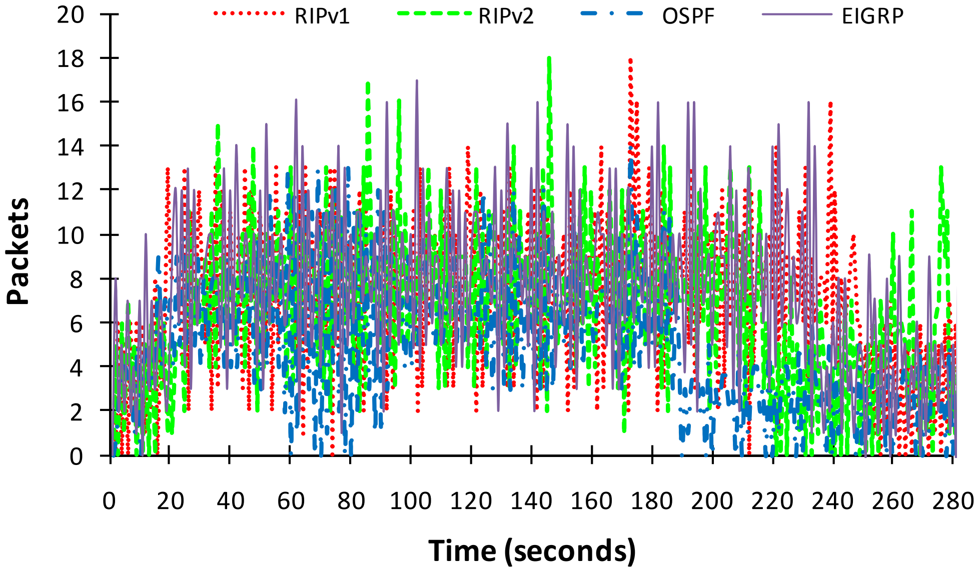

Figure 16. # of packets from PC_B to PC_A

6.5 Comparison of the # of bytes sent to the network

In this case we have measured the number of bytes that are travelling each second through the network when each routing protocol is used. The file transmitted from PC_A to PC_B and vice versa had a size of 1.36 GB. This information has been measured by Sniffer Pro in PC_C.

Fig. 17 shows the number of bytes per second when each routing protocol is running. The most unstable is OSPF because it has many peaks in its graph. During all the time, it is suffering oscillations between 7.000 and 11000 bytes. RIPv1 presents two main minimums. While RIPv1, OSPF and EIGRP reach around 10600 Bytes per second practically from the beginning, RIPv2 reached this value at around 50 seconds. On the other hand, while RIPv1, RIPv2 and OSPF decrease their value from 10600 to practically 0 at around 250 seconds, EIGRP does it at around 205 seconds. The average values have been: RIPv1=8983 Bytes per second, OSPF=7752 Bytes per second, EIGRP=7651 Bytes per second, and RIPv2=7559 Bytes per second.

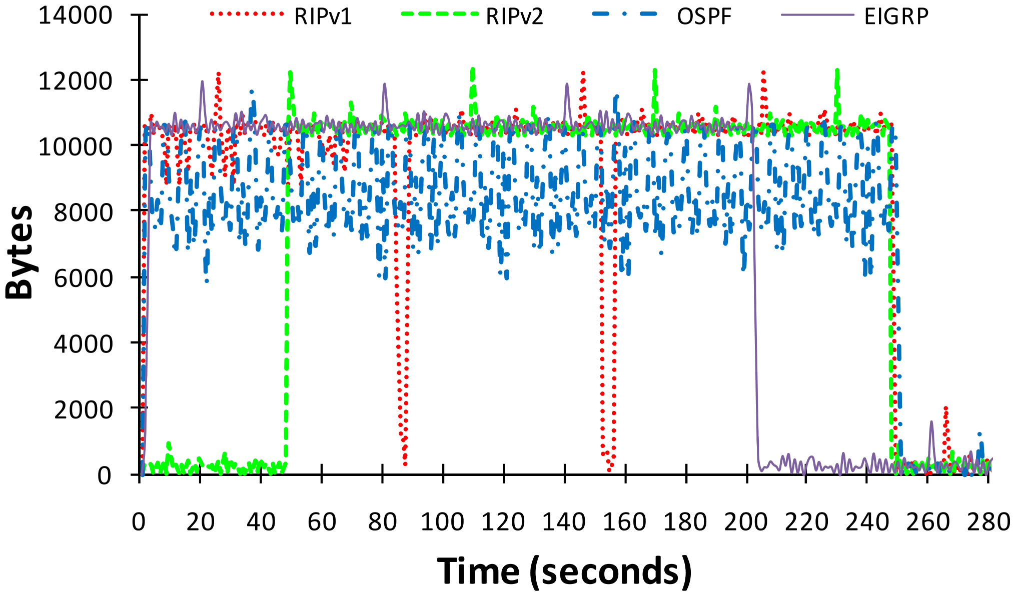

Figure 17. # of bytes from PC_A to PC_B

Fig. 18 compares the number of bytes per second sent from PC_B to PC_A when each routing protocol is running. In this case, the one that finishes first is OSPF (at around the second 180), the second one is RIPv2 (at around the second 205), and the last ones have been EIGRP and RIPv1. The average value of each one has been: RIPv1=5390 Bytes per second, EIGRP=4879 Bytes per second, RIPv2=4337 Bytes per second and the lowest one OSPF=3493 Bytes per second.

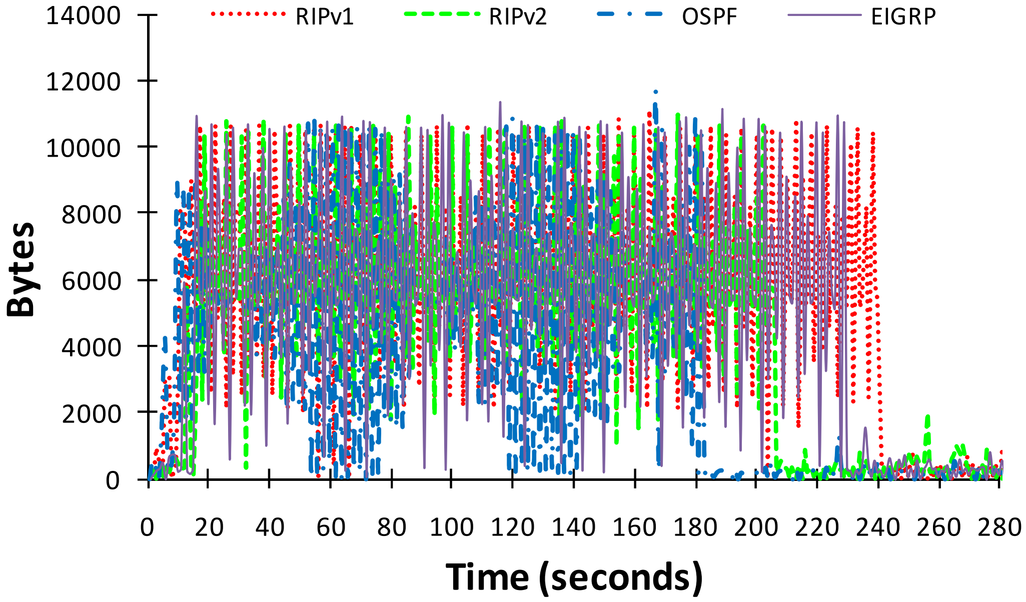

Figure 18. # of bytes from PC_B to PC_A

6.6 *Broadcast comparison*

Fig. 19 compares the broadcast traffic generated by each protocol in the network 192.168.99.0. We can see that the only protocol that generates broadcast packets is RIPv1. It makes sense because it is the only one that uses broadcast packets as update messages. We can observe that four messages are sent every 30 seconds.

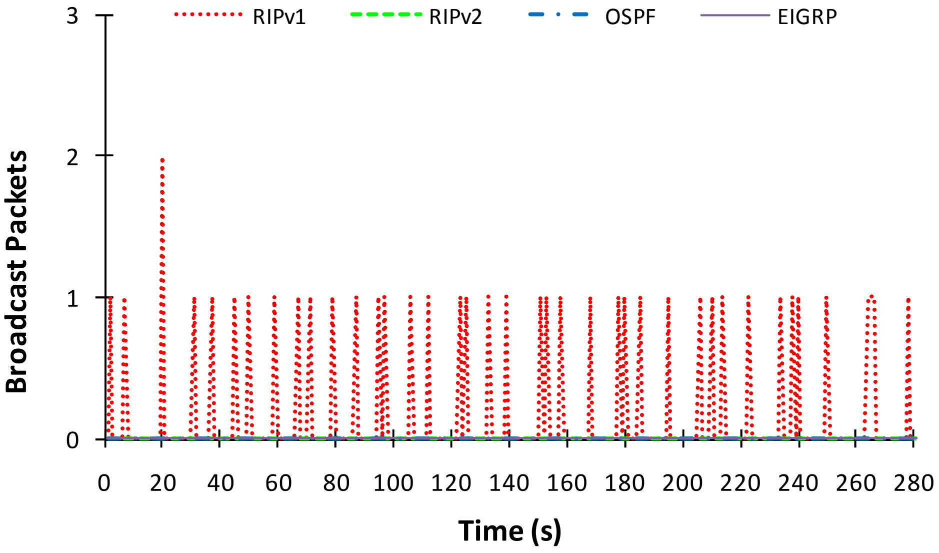

Figure 19. Broadcast packets comparison

7. Conclusion

In this paper, we have surveyed the most well known Interior Gateway Protocols (RIPv1, RIPv2, OSPF and EIGRP), which are used for communication between IP devices within a single administrative system. We have analyzed their main characteristics, like their packet format and their algorithm operation and their metrics. We have also compared their main features in order to know when and where they could be applied.

We have studied the behavior of the traffic when each routing protocol is running. Our measurements have been taken in two end PCs (PC_A and PC_B) and in the half of the path, where we located PC_C. All protocols had the same path in all cases in order to provide a better comparison.

We have observed that EIGRP is the protocol that first converges, followed by OSPF. The worst convergence case has been RIPv1. So, a network administrator should use EIGRP when a critical network is being administered because the network becomes stable fastest and the paths to the IP networks are found fastest when the network changes.

OSPF is the routing protocol with lowest average downlink and uplink bandwidth consumption and the one with most average downlink and uplink bandwidth consumption has been RIPv1. OSPF will be chosen by a network administrator when there are bandwidth constraints in some links of the network.

EIGRP is quite fast stabilizing when files are transferred through the network, but it has the highest average packets per second and its graph has many peaks. RIPv1 and OSPF are the ones that provide lowest average number of packets per second. Networks administrators should choose RIPv1 or OSPF when a fixed number of packets is reserved for the routing protocols and peaks of packets per second are not desired.

When we measured the number of bytes per second, we observed that the most unstable routing protocol was OSPF. In this case, RIPv1 has been the one with highest average number of bytes per second. So, RIPv1 should not be chosen when there are bandwidth limitations.

We have also observed that the only one that sends broadcast packets to the network has been RIPv1 (every 30 seconds).

In future works we will add more interior gateway routing protocols such as IS-IS and we will add some more experiments in order to have a complete view of their performance in other circumstances.

References

[1] International Organization for Standardization web site. Available at: http://www.iso.org/iso/home.html

[2] K. S. Siyan and T. Parker, TCP/IP Unleashet, Third Edition. Sams Publishing. 2004.

[3] CCIE Professional Development - Routing TCP-IP, Vol. I. Cisco Press. 1998.

[4] J. Hawkinson and T. Bates, Guidelines for creation, selection, and registration of an Autonomous System (AS). RFC - 1930. March 1996. Available at:

http://www.faqs.org/rfcs/rfc1930.html

[5] C. Hedrick, Routing Information Protocol. RFC – 1058. June 1988. Available at: http://www.faqs.org/rfcs/rfc1058.html

[6] G. Malkin, RIP Version 2. RFC – 2453. November 1998. Available at: http://www.faqs.org/rfcs/rfc2453.html

[7] J. Moy, OSPF Version 2. RFC – 2328. April 1998. Available at: http://www.faqs.org/rfcs/rfc2328.html

[8] Enhanced Interior Gateway Routing Protocol. Document ID: 16406. Available at Cisco web site: http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a0080094cb7.shtml

[9] Enhanced Interior Gateway Routing Protocol (EIGRP). Available at Cisco web site: http://www.cisco.com/en/US/tech/tk365/tk207/tsd_technology_support_sub-protocol_home.html

[10] EIGRP Packet Format. Available at Cisco web site: http://cisco.iphelp.ru/faq/5/ch06lev1sec6.html

[11] D. Pei, D. Massey, and L. Zhang, “Detection of Invalid Announcements in RIP protocols”. In procedings of IEEE Globecom 2003, 1-5 December (2003). San Francisco, California, USA.

[12] A. Shaikh, C. Isett, A. Greenberg, M. Roughan, and J. Gottlieb, “A Case Study of OSPF Behavior in a Large Enterprise Network”. In Procedings of Internet Measurement Workshop, Marseille (France). November 6-8, 2002.

[13] A. Basu and J. G. Riecke, “Stability issues in OSPF”. In Proceedings of ACM SIGCOMM, San Diego, CA, (USA). August 27-31, 2001.

[14] B. Albrightson, J. Garcia-Luna-Aceves, and J. Boyle, “EIGRP - a fast routing protocol based on distance vectors”. In Proceedings of NetWorld Interop Computer, Las Vegas, Nevada, (USA). May 2-6, 1994.

[15]A. Riesco and A. Verdejo, “The EIGRP Protocol in Maude”. Technical Report 3/07. Available at: http://eprints.ucm.es/6503/

[16]M. Nazrul and Md. A. Ullah, “Simulation-Based Comparative Study of EIGRP and OSPF for Real-Time Applications”, Master Thesis in Electrical Engineering Emphasis on Telecommunications, Blekinge Tekniska Hogskolan. Sweden, September, 2010. Available at: http://www.bth.se/fou/cuppsats.nsf/all/a3681538b6936d7fc125779e003f4b52/$file/Thesis_1053_Simulation-Based%20Comparative%20Study%20of%20EIGRP%20and%20OSPF%20for%20Real-Time%20Applications.pdf

[17] M. Nazrul and Md. A. Ullah, “Simulation Based EIGRP over OSPF Performance Analysis”, Master Thesis in Electrical Engineering Emphasis on Telecommunications, Blekinge Tekniska Hogskolan. Sweden, May, 2010. Available at: http://www.bth.se/com/mscee.nsf/attachments/4983_Thesis_Report_pdf/$file/4983_Thesis_Report.pdf

[18] OPNET website, http://www.opnet.com/

[19] E. S. Lemma, S. A. Hussain and W. W. Anjelo, “Performance Comparison of EIGRP/ IS-IS and OSPF/ IS-IS”, Master Thesis in Electrical Engineering Emphasis on Telecommunications. Blekinge Tekniska Hogskolan. Sweden, November 2009. Available at:

http://www.netlearning2002.org/fou/cuppsats.nsf/all/ac1cf74a90621ac4c12576ad003c45db/$file/Performance_Comparison_of%20EIGRPIS-IS_and_OSPFIS-IS_MEE09_77.pdf

[20] K. Al-Saud, H. Tahir, M. Saleh and M. Saleh. "Impact of MD5 Authentication on Routing Traffic for the Case of: EIGRP, RIPv2 and OSPF". Journal of Computer Science, Vol. 4, Issue 9, pp 721-728. 2008.

[21] C-C Chiang, C. Chen, D-L Jeng, S-J Wang and Y-K Ho, “The Performance and Security Evaluations of Internet Routing Protocols”. Journal of Informatics and Electronics, Vol.2, Issue.2, pp.21-27. March 2008.

[22] A. S. Tanenbaum, Computer Networks 4th edition, Prentice Hall Press, 2002.

[23] Diane Teare. Implementing Cisco IP Routing (ROUTE) Foundation Learning Guide, Cisco Press, Jun 28, 2010.

[24] Catalyst Switched Port Analyzer (SPAN) Configuration Example. Document ID: 10570. Available at Cisco web site: http://www.cisco.com/en/US/products/hw/switches/ps708/products_tech_note09186a008015c612.shtml

[25] Net Meter Website, at http://www.hootech.com/NetMeter/

[26] R. J. Shimonski, W. Eaton, U. Khan and Y. Gordienko, Sniffer Network Optimization and Troubleshooting Handbook, Syngress Publishing, Inc. 2002.

Copyright Disclaimer