# On the Convergence of the EIGRP Routing Protocol

Marjan Gusev

# On the Convergence of the EIGRP Routing Protocol

Dejan Spasov[1], Marjan Gusev[1], Sasko Ristov[1],

[1] Institute of Informatics, Faculty of Natural Science, Ss. Cyril and Methodius University,
Skopje, Macedonia
{dejan, marjan, sasko}@ii.edu.mk

**Abstract.** Distance vector routing protocols were popular choice among network engineers in the '80s due to their simplicity in implementation and low computational requirements. However, all distance vector protocols are prone to creation of routing loops; a phenomenon that is not desirable in modern computer networks. Cisco with their proprietary routing protocol - EIGRP claim that they solve the problem of creation of routing loops and improve the network convergence time. In this paper we give an overview of the EIGRP routing protocol with an emphasis on its convergence. We propose a mechanism that improves the convergence by taking into account the hop count that is reported by the neighboring routers.

**Keywords:** Routing protocols, convergence, RIP, OSPF, EIGRP.

## 1 Introduction

Modern networks have a huge impact on our communication, collaboration and interaction with other people in ways they never did before. We use computer networks in a variety of ways, including web applications, IP telephony, video conferencing, interactive gaming, electronic commerce, education, and more.

At the center of the network is the router. A router is a device that connects one network to another by determining the best path between them and reliably forwarding packets among them.

Routers have routing tables to keep the best paths for forwarding packets. Different protocols exist, so that the router can find the best path to remote networks.

### 1.1 Routing Protocols

Routers can learn about remote networks in one of the ways:

- Manually, configured by the administrator as static routes
- Automatically, supplied from a dynamic routing protocol

Static routes are commonly used when routing from a network to a stub network – network accessed by a single route.

Dynamic routing protocols are usually used in larger networks to ease the administrative and operational overhead of using only static routes. The dynamic routing protocols can be:

- Interior Gateway Protocols (IGP) used for routing inside an autonomous system
- Exterior Gateway Protocols (EGP) used for routing between autonomous systems

Typical IGP protocols are: RIP, IGRP[1], EIGRP, OSPF, and IS-IS. BGP is the only currently-viable EGP routing protocol.

IGP protocols can be classified as two types:

- Distance vector routing protocols
- Link-state routing protocols

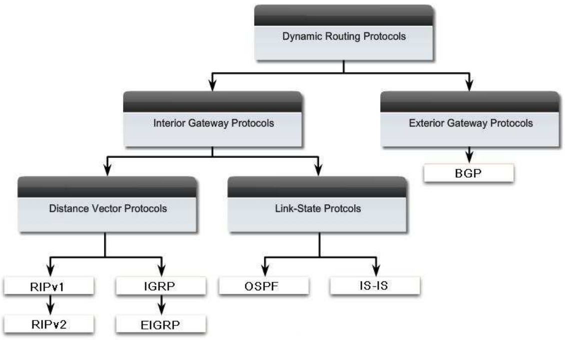

**Fig. 1.** Dynamic Routing Protocols Classification.

Routing protocols use a metric to determine the best path to a network. The metric used by the routing protocol RIP is hop count, which is the number of routers that a packet must traverse in reaching another network. OSPF uses bandwidth to determine the shortest path. EIGRP uses composite metric using Bandwidth, Delay, Reliability, Load to calculate preferred path to a network.

Distance vector means that routes shared as vector of distance and direction. Distance is defined in terms of metric and direction is simply the next hop router or exit interface.

Instead of distance vector routing protocols, routers using link-state routing protocol can create a complete network topology by gathering information from all routers in the network.

## 1.2 Classfull and Classless Routing Protocols

Historically, RFC1700 grouped the unicast ranges into specific sizes called class A, class B, and class C addresses. It also defined class D (multicast) and class E (experimental) addresses, as previously presented. In early days, networks were

---

[1] deprecated from 12.2 IOS and later

addressed as one of A, B or C unicast IP address class, with predefined subnet musk according the class belongs. This addressing is called classfull addressing.

Today networks use classless addressing. With the classless system, address blocks appropriate to the number of hosts in the network, without regard to the unicast class.

Routing protocols can be classified as classfull and classless according to networks addressing. Classful routing protocols do not send subnet mask information in routing updates. The routing protocols RIPV1 and IGRP are classful. A routing protocol did not need to include the subnet mask in the routing update because the network mask could be determined based on the first octet of the network address or IP Address class. Classful routing protocols cannot be used when a network is subnetted using more than one subnet mask, in other words classful routing protocols do not support variable length subnet masks (VLSM).

Opposite the classful, classless routing protocols include the subnet musk with the network address in routing updates. Classless routing protocols are required and used in most networks today because of their support for VLSM. Discontinuous networks are other example where classless routing protocol must be used.

Classless routing protocols are RIPv2, EIGRP, OSPF, IS-IS, BGP.

## 1.3 Convergence

Convergence is when routing tables on all routers are at a state of consistency. The network has converged when all routers have complete and accurate information about the network.

Convergence time is the time it takes routers to share information, calculate best paths, and update their routing tables. A network is not completely operable until the network has converged; therefore, most networks require short convergence times.

Convergence is both collaborative and independent. The routers share information with each other but every router must independently calculate the impacts of the topology change on their own routes, either a new router or interface is added, or a router or interface is down.

Convergence properties include the speed of propagation of routing information and the calculation of optimal paths. Routing protocols can be rated based on the speed to convergence; the faster the convergence, the better the routing protocol. Generally, RIP and IGRP are slow to converge, whereas EIGRP and OSPF are faster to converge.

## 1.4 EIGRP Routing Protocol

Enhanced Interior Gateway Routing Protocol (EIGRP) is a distance vector, classless routing protocol that was released in 1992 with IOS 9.21. It is Cisco proprietary protocol and only operates on Cisco routers.

It is a distance vector routing protocol using routing updates as vectors of distances outspreaded to directly connected neighbors.

Router configured with EIGRP routing protocol, send and receive messages to and from routers in the same Autonomous System (AS). EIGRP messages are encapsulated into a packet. It contains EIGRP header and data field called Type/Length/Value or TLV. Destination IP address is multicast 224.0.0.10, and the destination MAC address is also a multicast address: 01-00-5E-00-00-0A. (if the IP packet is encapsulated in an Ethernet frame)

**Fig. 2.** Encapsulated EIGRP Message.

Important fields in EIGRP header include the Opcode field and the Autonomous System Number field. Opcode specifies the EIGRP packet type, which can be:

- Update
- Query
- Reply
- Hello

EIGRP parameter TLV includes the weights that EIGRP uses for calculation its composite metric. Only bandwidth and delay are weighted by default. Both are equally weighted, therefore, the K1 field for bandwidth and the K3 field for delay are both set to 1. The other K values are set to zero by default.

There is also a hold time field, used for maximum time router should wait for the next hello packet. After hold time, router claims appropriate interface down.

IP Internal Routes TLV is used to advertise EIGRP routes within an autonomous system. Important fields for EIGRP protocol include:

- the metric fields (Delay and Bandwidth)
- the subnet mask field (Prefix Length)
- the Destination field (IP Address)

Delay is calculated as the sum of delays from source to destination in units of 10 microseconds. Bandwidth is the lowest configured bandwidth of any interface along the route.

The field “hop count” is very important for this paper, and will be discussed in Section 2.

EIGRP uses five different packet types.

- *Hello packets* are used to discover neighbors and to form adjacencies with those neighbors. EIGRP hello packets are multicasts and use unreliable delivery.
- *Update packets* are used to propagate routing information. Update packets are sent only when necessary. EIGRP updates contain only the routing information needed and are sent only to those routers that require it. EIGRP update packets use reliable delivery. Update packets are sent as a multicast when required by multiple routers, or as a unicast when required by only a single router.
- *Acknowledgement (ACK) packets* are sent when reliable delivery is used: update, query, and reply packets.
- *Query* and *reply packets* are used by DUAL when searching for networks and other tasks. Queries and replies use reliable delivery.

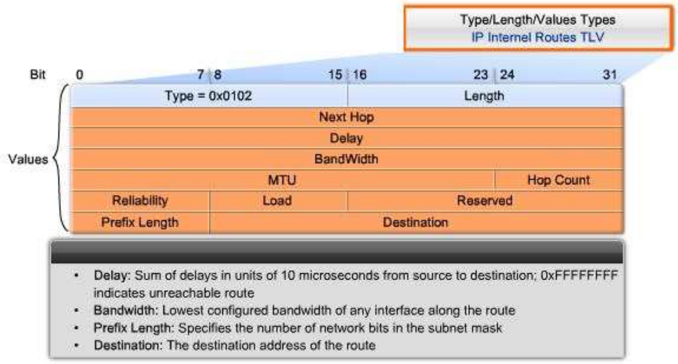

- Delay: Sum of delays in units of 10 microseconds from source to destination; 0xFFFFFFFF indicates unreachable route
- Bandwidth: Lowest configured bandwidth of any interface along the route
- Prefix Length: Specifies the number of network bits in the subnet mask
- Destination: The destination address of the route

**Fig. 3.** TLV: IP internal fields.

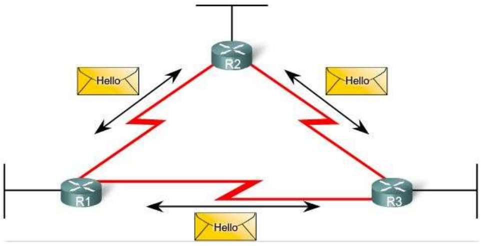

**Fig. 4.** EIGRP uses Hello packets to discover neighbors and to form adjacencies with them.

Hello packets are sent every 5 seconds on most networks. On multipoint nonbroadcast multi-access networks (NBMA), such as ATM, X.25 and Frame Relay, hello are sent on every 60 seconds. As long as router receives hello packets from a neighbor, the neighbor and routes in that direction remains valid, hold time is 3 times Hello interval, or 15 seconds on most networks, and 180 seconds on low speed NBMA networks.

EIGRP can be configured for authentication to encrypt and authenticate their routing information. It is good practice to authenticate transmitted routing information. This practice ensures that routers will only accept routing information from other routers that have been configured with the same password or authentication information. We must notice that authentication does not encrypt the router's routing table.

EIGRP uses composite metric to calculate the preferred path to the destination networks. Default composite metric is called the *bandwidth + delay* metric:

$$Metric = [K_1 * Bandwith + K_3 * Delay] . \tag{1}$$

Complete composite formula is:

$$Metric = \quad (2)$$
$$[K_1 * Bandwidth + (K_2 * Bandwidth) / (256 - Load) + K_3 * Delay] * .$$
$$*[K_5 / (\mathrm{Re}\,liability + K_4)]$$

If K2, K4, K5=0, then complete composite formula becomes default formula.

EIGRP uses an algorithm known as Diffusing Update Algorithm (DUAL) to determine the best path to the destination networks. EIGRP does not send periodic updates, but sends routing updates only when changes occurred such as a new link or unavailable link.

EIGRP's DUAL algorithm maintains a topology table separate from the routing table, which includes both the best path (with lowest metric) to a destination network and any backup paths that DUAL has determined to be loop-free.

These loop-free routes must meet a feasibility condition. Any backup path that meets this condition is guaranteed to be loop-free. Because EIGRP is a distance vector routing protocol, it is possible that there might be loop-free backup paths to a destination network that do not meet the feasibility condition. These paths are therefore not included in the topology table as a valid loop-free backup path by DUAL algorithm.

If a route becomes unavailable, DUAL will search its topology table for a valid backup path. If one exists, that route is immediately entered into the routing table. If one does not exist, DUAL performs a network discovery process to see if there happens to be a backup path that did not meet the requirement of the feasibility condition.

The main question that we try to answer is: can we improve the DUAL algorithm, so that every destination network has a valid backup path? Will this speed-up network convergence, reduce network discovery process and recalculations for metrics?

## 2 Improved Feasibility Condition for the EIGRP protocol

In a properly designed network, it is expected that multiple routes between any two destinations do exist. If one route becomes unavailable, the traffic flow must continue, if possible without interruptions, through a backup route. In this transient time, new best route must be chosen and propagated through the network. Long transient time has been considered as one of the main weaknesses of the distance vector routing protocols. For example, in the early days, one of the routing protocols - RIP v.1 propagated the alternative paths with the speed of 30 seconds per hop. In modern networks, with hundreds of routers, this is not an acceptable convergence time.

One solution to the problem, implemented in the EIGRP routing protocol, is to store in each router its neighbor’s routing tables. Then, when a route disappears from the routing table, the router will simply install a new route from the stored information. If a route is not found locally, the router will start time consuming diffusing computations to discover alternative routes.

In order to find the best loop-free paths to each network, EIGRP relies on the DUAL algorithm. Additional task for DUAL is finding the best back-up loop-free route, in case the primary route fails. In EIGRP, DUAL uses the well known bandwidth + delay metric for route comparison and for detecting routing loops. We will present a solution that can improve the convergence time by using additional metric for loop testing.

The best route to a particular network in the language of EIGRP is called *Feasible Distance* (*FD*). The next hop router, to which all packets for this network are forwarded, is called *Successor*. Each router sends its *FDs* to its neighbors. At the neighbor router, these *FDs* are called *Reported Distances* (*RD*). A router finds the *FD* and the successor to a particular network by minimizing over all collected *RDs* to which the cost to reach the neighbor is added. All rejected *RDs* participate in the process of finding the *Feasible Successor* (*FS*), i.e. the back-up router in case the successor fails.

In order to qualify for *FS*, first the neighbor router *i* must pass the *Feasibility Condition* (*FC*), i.e. the received reported distance $RD_i$ from *i* must be smaller than the Feasible Distance *FD*

$$FC:\ RD_i < FD. \tag{3}$$

Then, from all *i* that have passed the test (1), FS is chosen according to

$$FS = i\ \text{, such that } RD_i + c_i \text{ is minimal.} \tag{4}$$

where $c_i$ is the cost to get to the router *i*. Detailed proof that (3) is a sufficient condition for loop freedom can be found in [4].

Note that $RD_i$ and $c_i$ in (3) and (4) are measured in terms of bandwidth + delay. In addition, as a part of the EIGRP protocol, each neighbor reports the distance in hops to each network. Hop count is used to mark unreachable networks. However, we will use hop count to develop an improved Feasibility Condition criterion that allows more routers to become candidates for Feasible Successors.

Let $FD^{HC}$ denotes the Feasible Distance in hops to a network. Let $RD_i^{HC}$ is the reported distance in hops from the neighbor *i*. Then the route through the neighbor *i* is loop-free if at least one inequality in the *Improved Feasibility Condition* (i*FC*) holds true:

$$iFC:\ \begin{array}{c} RD_i < FD \\ RD_i^{HC} \le FD^{HC} \end{array}. \tag{5}$$

Then from all *i* that have passed (3), FS is chosen according to

$$FS = i\ \text{, such that } RD_i + c_i \text{ is minimal} \tag{6}$$

while $RD_i + c_i$ is still measured in bandwidth + delay. It is easy to see that $RD_i^{HC} \le FD^{HC}$ is also sufficient condition for securing loop freedom, since if a routing loop exist, then $RD_i^{HC}$ must be strictly greater than $FD^{HC}$. In general, if a neighbor router reports the distance to a network in *n* different metrics, then the feasibility condition can be made of *n* inequalities. In the case of EIGRP, using (5) and (6) is very convenient, since it does not require changes in the protocol.

In the next two examples we will demonstrate the advantage of (5) and (6) over (3) and (4). We will start with the example from the R. Graziani's lecture notes [1].

We assume that EIGRP is properly configured, and the network has converged. This means that all routers in the network have consistent routing information. We are interested in observing how the routers R1 and R2 find the Feasible Successor (FS) using (3) - (6). From R1's perspective, R3 is the Successor router, and the Feasible Distance is $FD_1 = 2172416$. From R2's perspective, R3 is, again, the Successor router, and the Feasible Distance is $FD_2 = 3014400$. R1 will receive Reported Distances $RD_2 = FD_2 = 3014400$ and $RD_2^{HC} = 1$, while the router R2 will receive $RD_1 = FD_1 = 2172416$ and $RD_1^{HC} = 1$.

At this point we assume that R1 and R2 have exchanged their RDs. If (3) is used as a FC test, then R2 will choose R1 as FS, since $FD_2 < RD_1$; but, R1 will have no FS since $FD_1 > RD_2$. On the other hand, if (5) is used as a FC, then R1 remains FS for R2, but now from the condition $RD_2^{HC} \le FD_1^{HC}$, R2 will become FS for R1 for the destination network. Thus, with figure 5 we show that more routers in a network are able to find Feasible Successors.

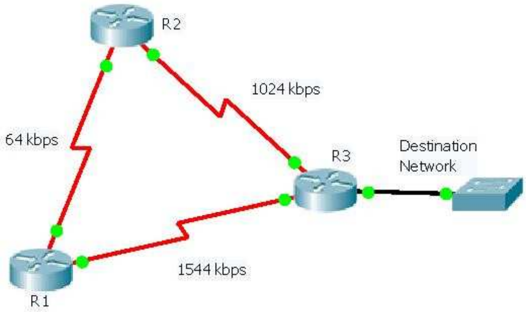

**Fig. 5.** Using the *iFC* more routers in the network are able to find back-up routes.

In the second example (fig. 6) we will focus on a single router and observe that with (5) better routes will be found.

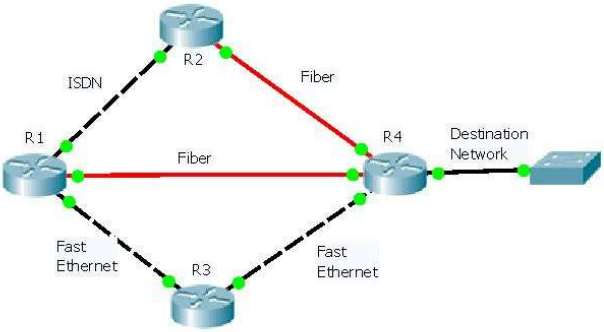

**Fig. 6.** Using the *iFC*, R1 has more routes to choose from and is able to find better back-up route through R3.

We will assume that the Fiber links have 1 Gbps and 10µs delay. Thus, R1 and R2 will have $FD_1 = FD_2 = 2816$, and R3 will have $FD_3 = 28160$. From R1's perspective, using (3), only R2 qualifies to be FS, since $FD_1 = RD_2 = 2816$. However, if we employ (5), then both R2 and R3 qualify as FS. The best route is then obtained with (6). The route R1-R2-R4 has cost that is mainly dictated by the slow ISDN link $c_1 = 40512256$, but the route R1-R3-R4 has cost $c_2 = 30720$. Since $c_2 < c_1$ the router R3 is chosen as FS. We conclude that the *iFC* helps individual routers to find better back-up routes.

## References

1. Graziani, R.: Lecture notes on CCNA and CCNP. Available at: http://www.cabrillo.edu/~rgraziani
2. Cisco Networking Academy: CCNA Exploration 4.0 - Routing Protocols and Concepts
3. Cisco Networking Academy: CCNP Advanced Routing Protocols
4. Garcia-Lunes-Aceves, J. J.: Loop-Free Routing Using Diffusing Computations. In: IEEE/ACM Transactions on Networking, vol. 1, no. 1, February 1993
5. Albrightson, B., Garcia-Lunes-Aceves, J. J., Boyle, J.: EIGRP – A Fast Routing Protocol Based on Distance Vectors [online].
6. Yee, J. R.: On the Internet Routing Protocol EIGRP: Is it Optimal? In: IFORS, 2005