Journal of Communication and Computer 9 (2012) 795-801

An Exploration of the Ambiguities of EIGRP Cost Matrix

Mahdi H Miraz[1]and Monir Bhuiyan[2]

1. Department of Computer Science & Software Engineering, University of Hail, Hail, Kingdom of Saudi Arabia

2. Institute of Information Technology, University of Dhaka, Dhaka, Bangladesh

Received: February 20, 2012 / Accepted: March 12, 2012 / Published: July 31, 2012.

Abstract: EIGRP (Enhanced Interior Gateway Routing Protocol), a Cisco proprietary product, is considered as one of the efficient routing algorithms currently being used in the field of networking and communication engineering. Like other routing algorithms, EIGRP uses its own Cost Matrix Calculation Formula to determine the best path to a destination. This formula consists of some constant values (“K” Values) upon which the path selection depends a lot. Although EIGRP is being used widely with efficient performance, it has been found that there are some ambiguities with its cost matrix. This paper is going to present a brief introduction to the background information of routing algorithms and explore the EIGRP cost matrix to find out those ambiguities lying behind it “K” Values.

Key words: Routing protocols, interior routing, cost matrix, optimization, EIGRP.

1. Introduction

Each and every available routing algorithm has its own way of calculating the cast known as “cost function” . This cost function differs from each other in different ways. The performance of a routing algorithm depends on how efficient is the cost function. Like others, Cisco’s EIGRP has its own cost function containing some K (constant) values in it.

With a brief introduction to the routing algorithms especially EIGRP, this paper is going to focus on the issues related to the ambiguity with the EIGRP metric.

1.1 Background: Category of Routing Protocols

Determining the best path to a given destination is one of the primary jobs of a router. A router gets the information about paths or routes from the static configuration entered by an administrator or dynamically from other routers, through routing protocols. So depending on this scenario, we get two types of routing.

Corresponding author: Mahdi H Miraz, lecturer, research fields: networking technologies, HCI, usability, cross-cultural IS issues. E-mail: m.miraz@uoh.edu.sa.

1.2 Static Routing

In static routing, an administrator manually initialize the router by defining routes to one or more destination networks and the routes do not change until s/he manually programs the changes Static routing is useful in networks that do not have multiple paths to any destination network. Static routing reduces the memory and processing burdens on a router.

1.3 Dynamic Routing

In dynamic routing, routers follow rules defined by routing protocols to exchange routing information and independently select the best path and the routes change automatically as neighboring routers update each other with new information.

Dynamic routing of TCP/IP can be implemented on a given network using one or more protocols. These protocols are often grouped according to where they are used. Routing protocols designed to work inside an autonomous system are categorized as IGPs (interior gateway protocols), and protocols that work between autonomous systems are classified as EGPs (exterior gateway protocols).

These protocols can be further categorized as either distance vector or link-state routing protocols, depending on their method of operation [1].

1.4 Link-State Routing

Link-state routing protocols collect route information from all other routers in the network or within a defined area of the network. Once all of the information is collected, each router calculates the best paths to all destinations in the network. The following are some link-state routing protocol functions:

(1) Respond quickly to network changes;

(2) Send triggered updates only when a network change has occurred;

(3) Send periodic updates known as link-state refreshes;

(4) Use a hello mechanism to determine the reachability of neighbors.

1.5 Distance Vector Routing

Distance vector routing protocols are based on the Bellman-Ford algorithm. Routers configured to use a distance vector routing protocol typically send their complete routing table at regular intervals to neighbour routers. Simple distance vector protocols, such as RIP and IGRP, broadcast their routing tables on all configured interfaces. This broadcast is also referred to as multicasting. Routers that use these protocols do not actually identify their neighbours for direct communication.

1.6 Hybrid Routing

Hybrid Routing is a combination of distance-vector routing, which works by sharing its knowledge of the entire network with its neighbors and link-state routing which works by having the routers tell every router on the network about its closest neighbors. Hybrid routing allows for rapid convergence but requires less processing power and memory as compared to link-state routing.

Although EIGRP (enhanced interior gateway routing protocol) is sometimes considered as a perfect example of a hybrid routing protocol [2], Cisco declines this claim and considers it to be a distance vector protocol [3].

2. How EIGRP Calculates Its Cost Function

EIGRP uses the minimum bandwidth on the path to a destination network and the total delay to compute routing metrics. Although it is possible to configure other metrics, Cisco does not recommend it, as it can cause routing loops in the network. The bandwidth and delay metrics are determined from values configured on the interfaces of routers in the path to the destination network [4].

EIGRP uses the same formula that IGRP uses to calculate its composite metric. However, to achieve a finer metric granularity, EIGRP calculates the total metric by scaling the bandwidth and delay metrics [5]. EIGRP uses the following formula to scale the bandwidth:

$$\text{bandwidth} = (10000000/\text{bandwidth}(i) * 256 \quad (1)$$

where bandwidth (i) is the least bandwidth of all outgoing interfaces on the route to the destination network represented in kilobits.

EIGRP uses the following formula to scale the delay:

$$\text{delay} = \text{delay}(i) * 256 \quad (2)$$

where delay(i) is the sum of the delays configured on the interfaces, on the route to the destination network, in tens of microseconds.

EIGRP uses these scaled values to determine the total metric to the network:

$$\text{metric} = [K1 * \text{bandwidth} + (K2 * \text{bandwidth}) / (256 - \text{load}) + K3 * \text{delay}] * [K5 / (\text{reliability} + K4)] \quad (3)$$

CISCO suggests that these K values should be used after careful planning. And it also warns us that mismatched k values prevent a neighbor relationship from being built, which can cause the network to fail to converge.

Let's consider K5 to be 0. So from Eq. (3) we get metric = [K1 * bandwidth + (K2 * bandwidth) / (256 - load) + K3 * delay] * [0 / (reliability + K4)] [Replacing K by 0]

or, metric = [k1 * bandwidth + (k2 * bandwidth)/ (256 - load) + k3 * delay].
[Using formula 0/x=0] (4)

Let's consider the other K values to be as follows:

- K1 = 1
- K2 = 0
- K3 = 1
- K4 = 0

These (including K5 = 0) are, in case known the default values of K defined by CISCO.

From Eq. (4) we get,

metric = [1 * bandwidth + (0 * bandwidth)/(256 - load) + 1 * delay] [Replacing the existing K values]

Or, Metric = bandwidth + delay [Using formulas 1) 1 * x = x and 0*x = 0] (5)

Combining Eqs. (1), (2) & (5) we get,

metric = (10000000/bandwidth(i)) * 256 + delay(i) * 256 (6)

This Eq. (5) or (6) is known as the simplified formula for default behaviour.

Cisco routers do not perform floating point math, so at each stage in the calculation, we need to round down to the nearest integer to properly calculate the metrics.

For instance Fig. 1, the author assume Router is going computing the best path to Network A. In this scenario, there are two ways: one through Router Four, with a minimum bandwidth of 56 and a total delay of 4400; and the other through Router Three, with a minimum bandwidth of 128 and a delay of 1200. Router One chooses the path with the lowest metric.

Let us compute the metrics using the equation we derived in the previous stage (simplified Eq. (6)).

Through Router Four is:

Given, minimum bandwidth = 56 k

total delay = 4000 + 200 + 200 = 4400

so metric = [(10000000/56) + 4400] x 256 = (178571 + 4400) x 256 = 182971 x 256 = 46840576

Through Router Three is:

Given, minimum bandwidth = 128 k

total delay = 200 + 200 + 2000 = 2400

So metric = [(10000000/128) + 2400] x 256 = (78125 + 2400) x 256 = 80525 x 256 = 20614400

So to reach Network A, Router One chooses the route through Router Three.

Here the bandwidth and delay values we used are those configured on the interface through which the router reaches its next hop to the destination network. For example, Router Two advertised Network A with

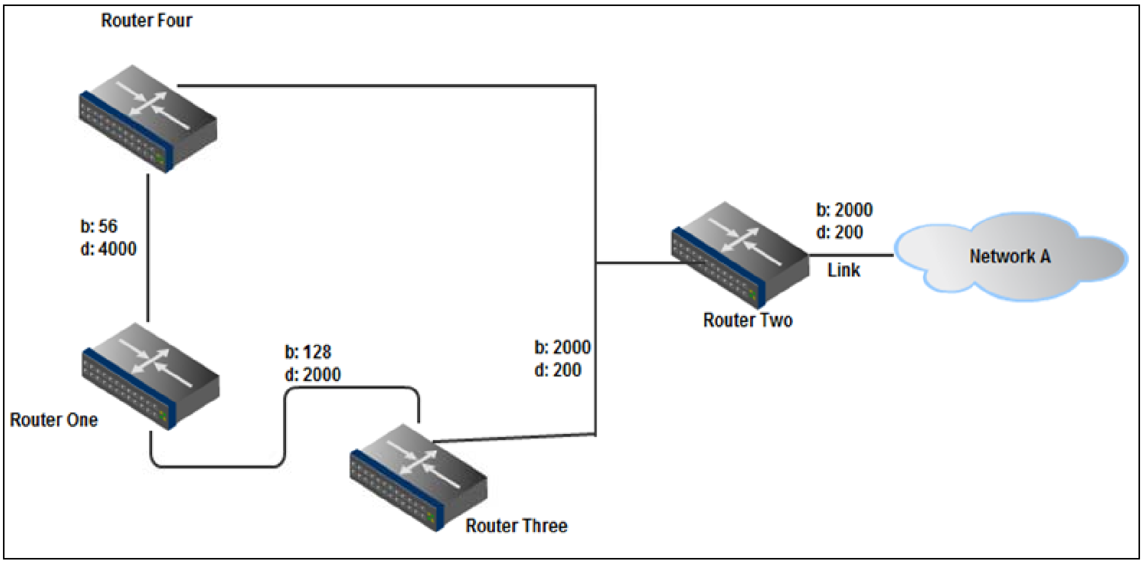

Fig. 1 Example of matrix calculation.

the delay configured on its Ethernet interface; Router Four added the delay configured on its Ethernet, and Router One added the delay configured on its serial.

3. Ambiguities of EIGRP Cost Matrix

The above equations look OK and seem to be very simple. But our research has found some ambiguities in the cost metric calculation function. Examples include:

EIGRP updates contain five metrics: minimum bandwidth, delay, load, reliability, and MTU (maximum transmission unit). Of these five metrics, by default, only minimum bandwidth and delay are used to compute best path. Unlike most metrics, minimum bandwidth is set to the minimum bandwidth of the entire path, and it does not reflect how many hops or low bandwidth links are in the path. Delay is a cumulative value which increases by the delay value of each segment in the path [6].

(1) All of these five metrics are recorded in the Topology table but why not in the Routing table? It is to mention that Topology table is only used when any existing route on the routing table is dead. So what is the point of considering MTU, load and reliability if only delay and minimum bandwidth is being used to calculate the Metric?

(2) Part of Eq. (3) is K5 / (reliability + K4). Now if K5 is considered 0 in Eq. (3), K4, for any value is automatically being omitted by the universal law of mathematics as 0/x = 0 and in our case x = reliability + K4. So what is the point of defining K4 = 0 as default values?

(3) Again, if K5 = 0, CISCO claims the simplified formula to be Eq. (6) But it misleads to an ambiguity. If K5 = 0, Eq. (3) derives metric = 0, which means all path's metrics are equal and (also vanishing.) In fact CISCO caims Eq. (3) applies if and only if K5>0 and otherwise it reduces to Eq. (4) [6-7].

(4) Delays are calculated in tens of milliseconds on interfaces but in Eq. (1) as tens of microseconds. This results is an additional scaling by a factor of 10 which increases the complexity of the calculation.

(5) Complexity has been also increased by calculating the delay and bandwidth as a factor of 256 (as shown in Eqs. (1) and (2)) where bandwidth(i) and delay(i) are either measure or configured values. The 256 arises from a storage difference (from IGRP to EIGRP) between 24 and 32 bits [8].

(6) Again, as Cisco routers do not perform floating point math, so at each stage in the calculation, we need to round down to the nearest integer to properly calculate the metrics. This forces to occur truncation

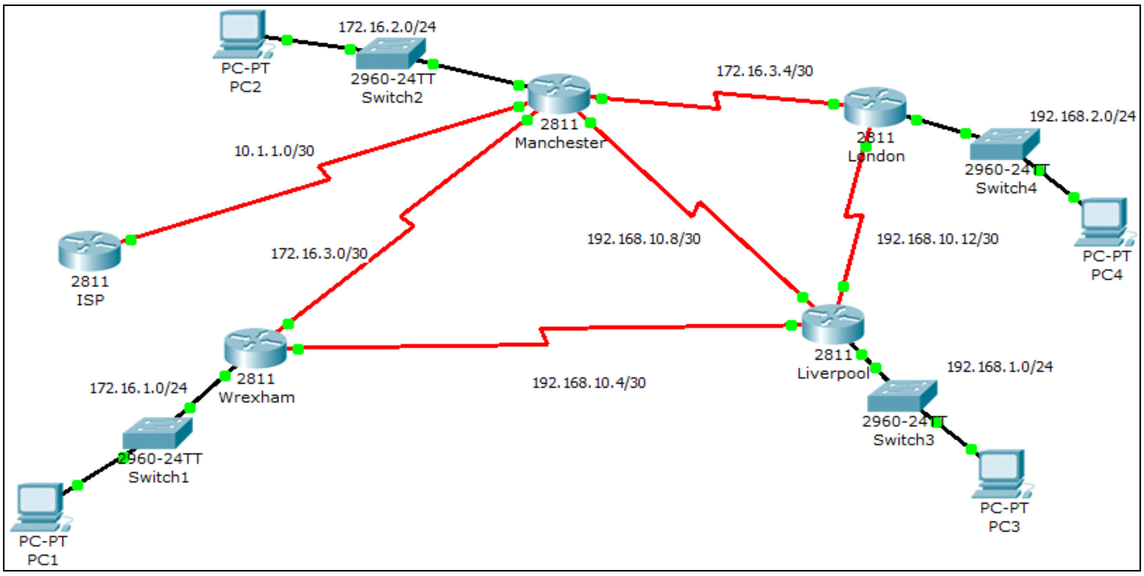

Fig. 2 Design of the simulated network.

and may lead to miscalculation of the lowest metric.

4. Simulation

Simulation of sample network, specially in academic research, is very important in computer networking and telecommunication. Not only it helps to get the insight view of a network, it also provides guidance for the future. Burbank [9] describes M&S (modeling and simulation) as a critical element in the design, development and T&E (test and evaluation) of any network product and solution. In this paper we have included the summarized results of some Cisco Packet Tracer 5.3.2 simulations while investigating the EIGRP behaviors for different "K" values.

We have configured EIGRP routing protocol in four routers namely Wrexham, Manchester, Liverpool and London. The Fig. 1 is a snapshot from Packet Tracer which demonstrate the whole network. The IP addresses and subnet masks of the respective interfaces of these four routers are provided in the flowing chart. The router "Manchester" has also been configured with loopback on the serial port (10.1.1.1) connecting to the ISP.

EIGRP requires all routers in the same routing domain to be configured with the same process ID. In our example, we have configured all the routers with 1 as the process ID. We also have configured TOS (type of service) with the default value of 1 as instructed by CISCO.

The composite metric formula used by EIGRP consists of 5 K values namely K1, K2, K3, K4 and K5. By default, K1 and K3 are set to 1, and K2, K4, and K5 are set to 0. As we have discussed earlier, the default behavior is that only the bandwidth and delay values are used in the computation of the default composite metric. We have configured the routers with different K values with default bandwidth to experiment what effects it have on the total routing process. The following table includes some of the results. Surprisingly enough, we have found that no changes in the cost metric takes place even if we include into the metric calculation by changing K4 = 1. The respective columns with K1 = 1, K2 = 1, K3 = 1, K4 = 0, K5 = 0 and K1 = 1, K2 = 1, K3 = 1, K4 = 1, K5 = 0 confirm this result.

Another interesting finding is that, when we assign k5 as 1, the neighbor adjacencies are lost and hence the routing table does not list any costs. We have further experimented the scenario by sending simple PDUs (with ICMP packets) from different sources to various destinations. A sample result of this scenario has been illustrated in the following figure. Only packets to the directly connected routers reached the destinations. Other routers/networks including directly connected networks on the FastEthernet interfaces were unreachable, and hence the packets were dropped.

Fig. 3 is a snapshoot of a packet tracer event list of the simulation using ICMP packets.

Table 1 IP addresses and subnet masks.

Device	Interface	IP Address	Subnet Mask
Wrexham	Fa0/0	172.16.1.1	255.255.255.0
	S0/0/0	172.16.3.1	255.255.255.252
	S0/0/1	192.168.10.5	255.255.255.252
Manchester	Fa0/0	172.16.2.1	255.255.255.0
	S0/0/0	172.16.3.2	255.255.255.252
	S0/0/1	192.168.10.9	255.255.255.252
	S0/1/0	10.1.1.1	255.255.255.252
	S0/1/1	172.168.3.6	255.255.255.252
Liverpool	Fa0/0	192.168.1.1	255.255.255.0
	S0/0/0	192.168.10.6	255.255.255.252
	S0/0/1	192.168.10.10	255.255.255.252
	S0/1/0	192.168.10.14	255.255.255.252
London	Fa0/0	192.168.2.1	255.255.255.0
	S0/0/0	192.168.10.13	255.255.255.252
	S0/0/1	192.168.10.5	255.255.255.252
PC1	NIC	172.16.1.10	255.255.255.0
PC2	NIC	172.16.2.10	255.255.255.0
PC3	NIC	192.168.1.10	255.255.255.0
PC4	NIC	192.168.2.10	255.255.255.0
ISP	S0/0/0	10.1.1.2	255.255.255.252

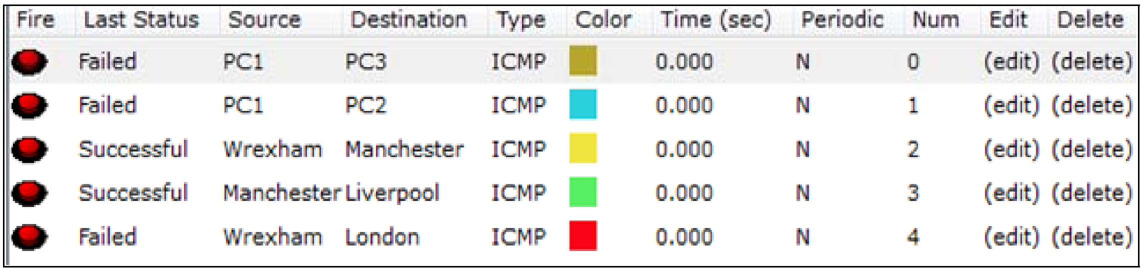

Fire	Last Status	Source	Destination	Type	Color	Time (sec)	Periodic	Num	Edit	Delete
	Failed	PC1	PC3	ICMP		0.000	N	0	(edit)	(delete)
	Failed	PC1	PC2	ICMP		0.000	N	1	(edit)	(delete)
	Successful	Wrexham	Manchester	ICMP		0.000	N	2	(edit)	(delete)
	Successful	Manchester	Liverpool	ICMP		0.000	N	3	(edit)	(delete)
	Failed	Wrexham	London	ICMP		0.000	N	4	(edit)	(delete)

Fig. 3 Events list of ICMP packets.

Table. 2 Simulation results for composite metrics.

Routers	Networks	Composite Metric for Different Combinations of the K Value			
		K1 = 1 K2 = 0 K3 = 1 K4 = 0 K5 = 0	K1 = 1 K2 = 1 K3 = 1 K4 = 0 K5 = 0	K1 = 1 K2 = 1 K3 = 1 K4 = 1 K5 = 0	K1 = 1 K2 = 1 K3 = 1 K4 = 1 K5 = 1
Wrexham	10.0.0.0/8	2297856	2304256	2304256	??
	172.16.2.0/24	2172416	2178816	2178816	??
	172.16.3.4/30	2681856	2688256	2688256	??
	192.168.1.0/24	2172416	2178816	2178816	??
	192.168.2.0/24 (via 192.168.10.6)	2684416	2690816	2690816	??
	192.168.2.0/24 (via 172.16.3.2)	2684416	2690816	2690816	??
	192.168.10.8/30	2681856	2688256	2688256	??
	192.168.10.12/30	2681856	2688256	2688256	??
Manchester	172.16.1.0/24	2172416	2178816	2178816	??
	192.168.1.0/24	2172416	2178816	2178816	??
	192.168.2.0/24	2172416	2178816	2178816	??
	192.168.10.4/30	2681856	2688256	2688256	??
	192.168.10.12/30	2681856	2688256	2688256	??
Liverpool	10.0.0.0/8	2297856	2304256	2304256	??
	172.16.0.0/16 (via 192.168.10.9)	2172416	2178816	2178816	??
	172.16.0.0/16 via 192.168.10.5	2172416	2178816	2178816	??
	192.168.2.0/24	2172416	2178816	2178816	??
London	10.0.0.0/8	2297856	2304256	2304256	??
	172.16.1.0/24	2684416	2690816	2690816	??
	172.16.2.0/24	2172416	2178816	2178816	??
	172.16.3.0/24	2681856	2688256	2688256	??
	192.168.1.0/24	2172416	2178816	2178816	??
	192.168.10.4/30	2681856	2688256	2688256	??
	192.168.10.8/30	2681856	2688256	2688256	??

5. Conclusions

As EIGRP is a hybrid of link-state and distance vector protocols, it is expected to be more efficient than those. But some features of EIGRP are inherited from its predecessor IGRP, so its delay and bandwidth calculations became more complicated and ambiguous. The above mentioned issues can be solved by deriving a properly mathematically functional new metric function which will be dependable on all the five factors (minimum bandwidth, delay, load, reliability, and MTU (maximum transmission unit)) that are recorded on the topology table. The issues stated in this paper should be taken into to serious consideration and a better cost metric for simple accurate calculation should be introduced which will work without any ambiguity.

References

[1] S. Keshav, An Engineering Approach to Computer Networking, Addison Wisley, 1999.

[2] Webopedia, Hybrid Routing Protocol, Home page, http://www.webopedia.com/TERM/H/hybrid_routing_protocol.html, 2011.

[3] Cisco, EIGRP: Chapter Introductionn, CCNA Explorer Resources, Cisco Network Academy Web site, http://curriculum.netacad.net/virtuoso/servlet/org.cli.delivery.rendering.servlet.CCServlet/LMS_ID=CNAMS,Theme=ccna3theme,Style=ccna3,Language=en,Version=1,RootID=knetlcms_exploration2_en_40,Engine=static/CHAPID=null/RLOID=null/RIOID=null/theme/cheetah.html?cid=0900000000&l1=en&l2=none&chapter=9, 2011.

[4] Cisco. Enhanced Interior Gateway Routing Protocol, Document ID: 16406 Web site, http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a0080094cb7.shtml ,2011.

[5] R. Graziani, A. Johnson, Routing Protocols and Concepts, CCNA Exploration Companion Guide, Cisco Press, 2008, pp. 263.

[6] Cisco, Setting a Preferred Route by Influencing EIGRP Metrics, Document ID: 13673 Web site, http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800c2d96.shtml, 2011.

[7] Cisco, IP Routing Protocols Commands: K through M Web site, http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/products_command_reference_chapter09186a008017d029.html#wp1040267, 2011.

[8] N. Houlden, V. Grout, J. McGinn, J. Davies, Extended end-to-end cost metrics for improved dynamic route calculation, in: Proceedings of the 6th International Network Conference (INC 2006), University of Plymouth, 2006, pp. 89-96.

[9] Burbank, Modeling and simulation: A practical guide for network designers and developers, IEEE Communications Magazine 47 (2009) 118.