Maximizable Routing Metrics

Mohamed G. Gouda
Department of Computer Sciences
The University of Texas at Austin
Austin, Texas 78712-1188, USA
gouda@cs.utexas.edu

Marco Schneider
SBC Technology Resources Inc.
9505 Arboretum Blvd.
Austin, Texas 78759-7260, USA
mschneider@tri.sbc.com

Abstract

We develop a theory for deciding, for any routing metric and any network, whether the messages in this network can be routed along paths whose metric values are maximum. In order for the messages in a network to be routed along paths whose metric values are maximum, the network needs to have a rooted spanning tree that is maximal with respect to the routing metric. We identify two important properties of routing metrics: boundedness and monotonicity, and show that these two properties are both necessary and sufficient to ensure that any network has a maximal tree with respect to any (bounded and monotonic) metric. We also discuss how to combine two (or more) routing metrics into a single composite metric such that if the original metrics are bounded and monotonic, then the composite metric is bounded and monotonic. Finally we show that the composite routing metrics used in IGRP (Inter-Gateway Routing Protocol) and EIGRP (Enhanced IGRP) are bounded but not monotonic.

1. Introduction

The primary task of a routing protocol is to generate and maintain a tree with the appropriate desirable properties (e.g., shortest path, maximum bandwidth, etc.). As the topology and available resources of a network change over time, this may necessitate that the tree be updated or rebuilt so that the resulting tree is maximal with respect to a given routing metric. Within the Internet, the common basis for IP routing is the *shortest path tree*. Conventional routing protocols such as OSPF and RIP utilize link-state and distance-vector routing in order to build a shortest path tree that minimizes the latency encountered when routing a datagram from one location to another. The explosion of traffic in the Internet due to the World Wide Web, along with the demand for multimedia applications such as real-time audio and videoconferencing, is necessitating its rearchitecture so as to provide more flexible quality of service types that take into account bandwidth and other network properties in addition to latency.

In recent years several authors have investigated alternative routing metrics based on bandwidth and other measures besides distance. For instance [GS94, GS95, Sch97] introduced the notion of maximum flow trees and provided a stabilizing distance-vector based protocol for their construction. Define the flow of a path as the minimum capacity of an edge along that path. A maximum flow tree in a network is a rooted spanning tree of the network wherein the path from any node to the root is a maximum flow path. Maximum flow trees were also independently introduced and studied by [WC95] in the context of the distance-vector paradigm. See also [CD94] wherein a metric is provided for depth first search tree construction.

In this paper we develop a theory for deciding, for any routing metric and any network, whether the messages in this network can be routed along paths whose metric values are maximum. In order for the messages in a network to be routed along paths whose metric values are maximum, the network needs to have a rooted spanning tree that is maximal with respect to the routing metric. We identify two important properties of routing metrics: boundedness and monotonicity, and show that these two properties are both necessary and sufficient to ensure that any network has a maximal tree with respect to any (bounded and monotonic) metric. Examples of trees based upon bounded and monotonic routing metrics include shortest path trees (distance-vector), depth first search trees, maximum flow trees and reliability trees. We also discuss how to combine two (or more) routing metrics into a single composite metric such that if the original metrics are bounded and monotonic, then the composite metric is bounded and monotonic. Finally we argue that the composite routing metrics used in IGRP (Inter-Gateway Routing Protocol) and EIGRP (Enhanced IGRP) are bounded but not monotonic.

The rest of this paper is organized as follows. First in Section 2 we define the notion of a routing metric and give examples. Next in Section 3 we provide two properties of routing metrics, boundedness and monotonicity, which we show to be necessary and sufficient for any network to have a spanning tree that is maximal with respect to a given routing metric. Then in Section 4 we show how to compose routing metrics such that maximality is preserved. Next in Section 5 we analyze the IGRP/EIGRP protocols and show that their metrics are bounded but not monotonic. Finally we make some concluding remarks in Section 6.

2. Routing Metrics

A network is an undirected graph where each node represents a computer, and each undirected edge between two nodes represents a communication channel between the two computers represented by the nodes.

To simplify our discussion of routing data messages in a network N, we assume that all data messages that are generated in all the nodes of N are to be routed to a distinct node in N. This node is called the *root* of N. It is straightforward to extend our discussion to the case where data messages are routed to arbitrary nodes.

Let N be a network and r be the root of N. In order to route the data messages generated at all the nodes of N to node r, a spanning tree rooted at node r, is maintained in N. When a node generates (or receives from one of its neighbors) a data message, the node forwards the data message to its parent in the spanning tree. Each data message is forwarded from any node to the node's parent in the tree until the data message reaches the root of the tree, node r.

In any network N with root r, there are many spanning trees whose root is r. Thus, our goal is to find a spanning tree that is maximal with respect to a given "routing metric". Consider for example a metric based upon flow. The flow of each node in a flow tree is computed by applying the *min* function to the flow of its parent in the tree and the capacity of the edge to its parent in the tree. In a maximum flow tree, for every node, its path along the tree has the maximum possible flow of any path to the root. Figure 4 contains an example of a maximum flow tree which will be explained in more detail later. Maximum flow trees are useful for the routing of virtual circuits. Our discussion can also be extended to achieve the alternative goal of finding a spanning tree that is minimal with respect to a given routing metric. Consider for example a metric based upon distance. The distance of each node in a minimum distance tree is computed by applying the addition function to the distance of its parent in the tree and the cost of the edge to its parent in the tree. In a minimum distance tree, for every node, its path along the tree has the minimum possible distance of any path to the root. The minimum distance tree is more commonly known as the shortest path tree. Figure 5 contains an example of a shortest path tree that will be explained in more detail later.

Before we give a formal definition of the general concept of routing metrics, we illustrate this concept by the following example. Consider the network in Figure 1. This network has six edges and five nodes, and the root of the network is the node labeled 0. Associated with each edge $\{i, j\}$ in this network is a weight w_{ij}. This network has many spanning trees whose root is node 0; two of those spanning trees T_1 and T_2 are shown in Figures 2 and 3 respectively. In each of these spanning trees, the metric value m_i of each node i can be computed as follows, where mr is a metric value, and *met* is a function that takes a metric value and an edge weight as inputs and computes a metric value.

1. If node i is the root, then $m_i = mr$.
2. If node i is not the root and node j is the parent of node i in the tree and w_{ij} is the weight of edge $\{i, j\}$, then $m_i = met(m_j, w_{ij})$.

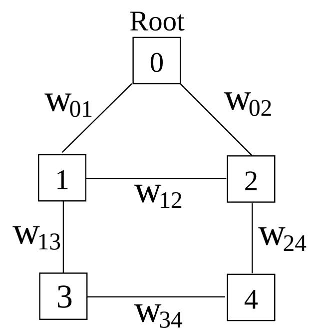

Figure 1. Network N

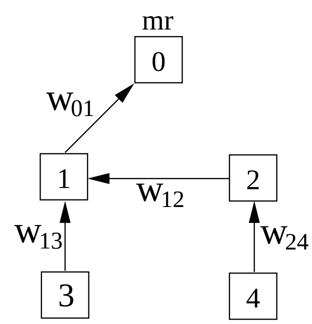

Figure 2. Spanning Tree T_1

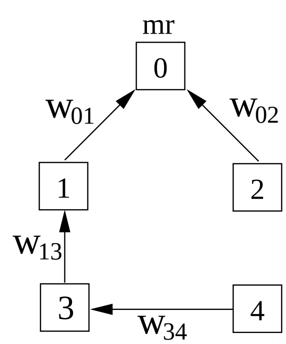

Figure 3. Spanning Tree T_2

For example the metric values of all nodes in the spanning tree T_1 in Figure 2 are as follows:

$$\begin{aligned} m_0 &= mr \\ m_1 &= met(m_0, w_{01}) \\ m_2 &= met(m_1, w_{12}) \\ m_3 &= met(m_1, w_{13}) \\ m_4 &= met(m_2, w_{24}) \end{aligned}$$

Also the metric values of all nodes in the spanning tree T_2 in Figure 3 are as follows:

$$\begin{aligned} m_0 &= mr \\ m_1 &= met(m_0, w_{01}) \\ m_2 &= met(m_1, w_{02}) \\ m_3 &= met(m_1, w_{13}) \\ m_4 &= met(m_3, w_{34}) \end{aligned}$$

The goal of a routing algorithm is to construct a tree that simultaneously maximizes the metric values of all of the nodes with respect to some ordering $\prec$. With this in mind we are now ready to formally define the concept of a routing metric.

Routing Metric

A routing metric for a network N is a six-tuple $(W, wf, M, mr, met, \prec)$ where:

1. W is a set of edge weights
2. wf is a function that assigns to each edge $\{i, j\}$, in N, a weight $wf(\{i, j\})$ in W
3. M is a set of metric values
4. mr is a metric value in M that is assigned to the root of network N
5. *met* is a metric function whose domain is $M \times W$ and whose range is M (it takes a metric value and an edge value and returns a metric value)
6. $\prec$ is a binary relation over M, the set of metric values, that satifies the following four conditions for arbitrary metric values m, m', and m'' in M:
   (a) Irreflexivity: $m \neg \prec m$
   (b) Antisymmetry: if $m \prec m'$ then $m' \not\prec m$
   (c) Transitivity: if $m \prec m'$ and $m' \prec m''$ then $m \prec m''$
   (d) Totality: $m \prec m'$ or $m' \prec m$ or $m = m'$

   Notice that the less-than relation "$<$" over the integers satisfies these four conditions.

We also require that every metric value $m \in M$ satisfies the following *utility condition*: For any metric value $m \in M$ there is a non-empty sequence of edge weights $w_1, w_2, \ldots, w_{k+1}$ ($w_i \in W$) and a sequence metric values $m_1, m_2, \ldots, m_{k+1}$ ($m_i \in M$) such that the following holds:

$$\begin{aligned} m_1 &= mr \\ m_2 &= met(m_1, w_1) \\ &\cdot \\ &\cdot \\ &\cdot \\ m_{k+1} &= met(m_k, w_k) \\ m &= m_{k+1} \end{aligned}$$

If there is a metric value in M that does not satisfy the utility condition, then this value can be removed from M and never missed.

We present two examples of metrics, the flow metric and the distance metric.

Flow Metric

The *flow metric* $(W, wf, M, mr, min, \prec)$ is defined as follows:

1. W is a subset of the non-negative integers which make up the set of possible edge capacities of the network
2. wf assigns each edge a capacity
3. M is a subset of the non-negative integers which make up the set of possible flow (metric) values
4. mr is chosen to be the maximum edge capacity that appears in the network
5. *min* is simply the minimum function which returns the minimum of two non-negative integers
6. $\prec$ is the less-than relation over the non-negative integers.

Distance Metric

The *distance metric* $(W, wf, M, mr, plus, \prec)$ is defined as follows:

1. W is a subset of the non-negative integers which make up the set of possible edge costs of the network
2. wf assigns to each edge a cost
3. M is a subset of the non-negative integers which make up the set of possible distance (metric) values
4. mr is equal to zero, the distance of the root from itself
5. *plus* is the addition function which returns the sum of two non-negative integers
6. $\prec$ is the greater-than relation over the non-negative integers.

Note that we use greater than instead of less-than for distance so that when we "maximize", we are really minimizing.

Maximal Tree

Let N be a network with root r, and let $(W, wf, M, mr, met, \prec)$ be a routing metric for N. A spanning tree T of N is called a *maximal tree* with respect to this routing metric iff for every spanning tree T' and every node i in N,

$$(m'_i \prec m_i) \ \vee \ (m'_i = m_i)$$

where m'_i is the metric value of node i in tree T', and m_i is the metric value of node i in tree T.

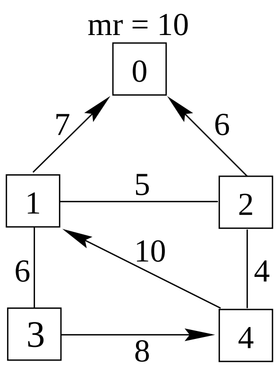

Figure 4. Maximum Flow Tree

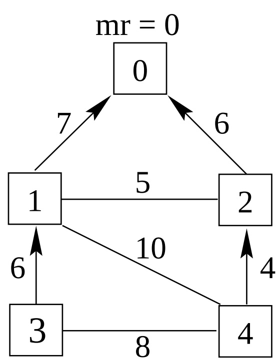

Figure 5. Shortest Path Tree

Consider the network in Figure 4 where each edge is labeled with an integer capacity. The maximum flow tree of this network consists of those edges which have been directed. This tree is maximal with respect to the flow metric. For each of the nodes in the network, its maximum possible flow value is obtained in its path along the overlayed tree. Node 1 has a flow 7 via node 0, node 2 has a flow of 6 via node 0, node 4 has a flow of 7 via node 1, and node 3 has a flow of 7 via node 4. Now consider the network in Figure 5 where each edge is labeled with an integer cost. The shortest path tree of this network consists of those edges which have been directed. This tree is maximal with respect to the distance metric as defined earlier. For each of the nodes in the network, its minimum possible distance value is obtained in its path along the overlayed tree. Node 1 has a distance of 7 via node 0, node 2 has a distance of 6 via node 0, node 3 has a distance of 13 via node 1, and node 4 has a distance of 10 via node 2.

Note that it is not possible to simultaneously maximize the distance of every node in a tree. Consider that in order to provide any node with its longest possible path we will be required to place other nodes along its path to the root and thus deprive them of their maximum values.

3. Properties of Maximizable Routing Metrics

In this section we identify two important properties of routing metrics, namely boundedness and monotonicity. It turns out that these two properties are both necessary and sufficient for constructing a maximal tree with respect to any routing metric. For detailed the proofs of the theorems in this section see [Sch97].

Boundedness: A routing metric $(W, wf, M, mr, met, \prec)$ is *bounded* iff the following condition holds for every edge weight w in W, and every metric value m in M:

$$met(m, w) \prec m \ \vee met(m, w) = m$$

Monotonicity: A routing metric $(W, wf, M, mr, met, \prec)$ is *monotonic* iff the following condition holds for every edge weight w in W, and every pair of metric values m and m' in M:

$$m \prec m' \ \Rightarrow \ (met(m, w) \prec met(m', w) \ \vee \ met(m, w) = met(m', w))$$

Theorem 3.1 *(Necessity of Boundedness)*
If a routing metric is chosen for any network N, and if N has a maximal spanning tree with respect to the metric, then the routing metric is bounded.

Theorem 3.2 *(Necessity of Monotonicity)*
If a routing metric is chosen for any network N, and if N has a maximal spanning tree with respect to the metric, then the routing metric is monotonic.

Theorem 3.3 *(Sufficiency of Boundedness and Monotonicity)*
If a routing metric is chosen for any network N, and if this routing metric is both bounded and monotonic, then N has a maximal spanning tree with respect to this metric.

It is easy to verify that both the flow metric (with $\prec$ defined as $<$) and the distance metric (with $\prec$ defined as $>$) are bounded and monotonic.

4. Composition of Maximizable Routing Metrics

In Section 2 we gave a formal definition of a routing metric and gave two examples based on flow and distance. In this section we look at composite routing metrics. In particular we address the following question. Given two bounded and monotonic routing metrics:

1. $(W_1, wf_1, M_1, mr_1, met_1, \prec_1)$
2. $(W_2, wf_2, M_2, mr_2, met_2, \prec_2)$

how do we combine these two metrics into a single metric $(W, wf, M, mr, met, \prec)$ that is both bounded and monotonic.

Clearly the combined metric needs to satisfy the following conditions:

$$\begin{aligned} W &= W_1 \times W_2 \\ wf &= \langle wf_1, wf_2 \rangle \\ M &= M_1 \times M_2 \\ mr &= \langle mr_1, mr_2 \rangle \\ met &= \langle met_1, met_2 \rangle \end{aligned}$$

According to these conditions, each edge weight in the combined metric is a pair $\langle w_1, w_2 \rangle$ where w_1 is an edge weight in the first metric and w_2 is an edge weight in the second metric. Also each value of the combined metric is a pair $\langle m_1, m_2 \rangle$, where m_1 is a value of the first metric and m_2 is a value of the second metric.

These conditions define W, wf, M, mr, and *met* of the combined metric, but they do not define the relation $\prec$ of the combined metric. This relation needs to be defined carefully to ensure that it satifies the four conditions of *irreflexivity*, *antisymmetry*, *transitivity*, and *general totality* (defined in Section 2).

Next we give a possible definition of relation $\prec$ in terms of the two relations $\prec_1$ and $\prec_2$. Let $\prec_1$ be a relation over a set M_1, and $\prec_2$ be a relation over a set M_2. A relation $\prec$ over the set $M_1 \times M_2$ is called a $\langle \prec_1, \prec_2 \rangle$-*sequence* iff the following condition holds:

For every $\langle m_1, m_2 \rangle$ and $\langle n_1, n_2 \rangle$ in $M_1 \times M_2$, $\langle m_1, m_2 \rangle \prec \langle n_1, n_2 \rangle$ iff either $m_1 \prec n_1$ or $m_1 = n_1$ and $m_2 \prec n_2$.

The intuition behind the $\langle \prec_1, \prec_2 \rangle$-sequence relation is that when comparing two composite metric values we prefer the one with the larger first indice. If the first two indices are the same, we prefer the one with the larger second indice. It is straightforward to show that if two relations $\prec_1$ and $\prec_2$ satisfy the four conditions of irreflexivity, antisymmetry, transitivity, and totality, then the $\langle \prec_1, \prec_2 \rangle$-sequence relation satisfies the same four conditions.

Theorem 4.1 *(Boundedness of the Sequenced Metric)*
If the following two routing metrics are bounded:

1. $(W_1, wf_1, M_1, mr_1, met_1, \prec_1)$
2. $(W_2, wf_2, M_2, mr_2, met_2, \prec_2)$

then the following composite metric is bounded where $\prec$ is the $\langle \prec_1, \prec_2 \rangle$-sequence relation:

$$(W_1 \times W_2, \langle wf_1, wf_2 \rangle, M_1 \times M_2, \langle mr_1, mr_2 \rangle, \langle met_1, met_2 \rangle, \prec)$$

Strict Monotonicity

A routing metric $(W, wf, M, mr, met, \prec)$ is called *strictly monotonic* iff the following condition holds for every edge weight w in W and every pair of metric values m and m' in M:

$$m \prec m' \Rightarrow met(m, w) \prec met(m', w)$$

Theorem 4.2 *(Monotonicity of the Sequenced Metric)*
If the following two routing metrics are monotonic:

1. $(W_1, wf_1, M_1, mr_1, met_1, \prec_1)$
2. $(W_2, wf_2, M_2, mr_2, met_2, \prec_2)$

and if the first metric is also strictly monotonic then the following composite metric is monotonic where $\prec$ is the $\langle \prec_1, \prec_2 \rangle$-sequence relation:

$$(W_1 \times W_2, \langle wf_1, wf_2 \rangle, M_1 \times M_2, \langle mr_1, mr_2 \rangle, \langle met_1, met_2 \rangle, \prec)$$

It is easy to see that while the distance metric is strictly monotonic, the flow metric is not. Thus the sequence metric formed from distance and then flow is both bounded and monotonic. However it can easily be shown that the sequence metric formed from flow and then distance is bounded but not monotonic and thus there is not a maximal tree with respect to this metric [Sch97].

As another example of a strictly monotonic metric we introduce the reliability metric.

Reliability Metric

The *reliability metric* $(W, wf, M, mr, times, \prec)$ is defined as follows:

1. W is a subset of the real numbers p such that $0 \leq p \leq 1$
2. wf assigns each edge a reliability
3. M is a subset of the real numbers p such that $0 \leq p \leq 1$ which make up the set of possible reliability (metric) values
4. mr is equal to 1
5. *times* is the multiplication function over real numbers p such that $0 \leq p \leq 1$
6. $\prec$ is the less-than relation over real numbers p such that $0 \leq p \leq 1$

The reliability of a path is a measure of how likely it is to either corrupt or drop data and is the product of the reliabilities of the edges along it. Because the reliablility metric is bounded and strictly monotonic it may be sequenced with the distance metric or the flow metric while preserving boundedness and monotonicity.

We conclude with a special case of metric composition that preserves boundedness but not necessarily monotonicity. First we need to slightly generalize our definition of a routing metric from Section 2. Recall that in this definition, the relation $\prec$ is required to satisfy (among other conditions) the following *totality* condition.

Totality:
For every pair of metric values m and m' in M, $m \prec m' \vee m' \prec m \vee m = m'$.

We slightly generalize this condition as follows:

Generality Totality:
There is an equivalence relation $\equiv$ over M, such that for every pair of metric values m and m' in M, $m \prec m' \vee m' \prec m \vee m \equiv m'$.

The general totality condition reduces to the totality condition by choosing the equivalence relation $\equiv$ to be the equality relation $=$.

Consider the following two routing metrics:

1. $(W_1, wf_1, M_1, mr_1, met_1, \prec_1)$
2. $(W_2, wf_2, M_2, mr_2, met_2, \prec_2)$

where each of M_1 and M_2 is the set of all integers and each of $\prec_1$ and $\prec_2$ is the less-than relation over integers.

These two metrics can be combined into the following composite metric which we call an *Additive Integer Metric*:

$$(W_1 \times W_2, \langle wf_1, wf_2\rangle, M_1 \times M_2, \langle mr_1, mr_2\rangle, \langle met_1, met_2\rangle, \prec)$$

where $\prec$ is defined as follows (with "$+$" as the integer addition operator, and "$<$" as the less-than relation over integers):

$$\langle m_1, m_2\rangle \prec \langle m_1', m_2'\rangle \Leftrightarrow (m_1 + m_2) < (m_1' + m_2')$$

Notice that $\prec$ satisfies the general totality condition by defining the equivalence relation $\equiv$ as follows:

$$\langle m_1, m_2\rangle \equiv \langle m_1', m_2'\rangle \Leftrightarrow (m_1 + m_2) = (m_1' + m_2')$$

Theorem 4.3 *(Boundedness of the Additive Integer Metric) If the following two routing metrics are bounded:*

1. $(W_1, wf_1, M_1, mr_1, met_1, \prec_1)$

2. $(W_2, wf_2, M_2, mr_2, met_2, \prec_2)$

where each of M_1 and M_2 is the set of all integers and each of $\prec_1$ and $\prec_2$ is the less-than relation over integers. Then the following metric is bounded:

$$(W_1 \times W_2, \langle wf_1, wf_2\rangle, M_1 \times M_2, \langle mr_1, mr_2\rangle, \langle met_1, met_2\rangle, \prec)$$

where $\prec$ is defined as follows (with "$+$" as the integer addition operator, and "$<$" as the less-than relation over integers):

$$\langle m_1, m_2\rangle \prec \langle m_1', m_2'\rangle \Leftrightarrow (m_1 + m_2) < (m_1' + m_2')$$

The above theorem generalizes to the composition of multiple metrics defined over the integers. In the following section we look at a particular instance of additive integer metric composition that corresponds to the well known IGRP/EIGRP protocols. In particular we look at an additive integer metric composed from inverse bandwidth and distance and we show that it is not monotonic.

5. Analysis of IGRP

Both the Inter-Gateway Routing Protocol (or IGRP for short) and the Enhanced Inter-Gateway Routing Protocol (EIGRP for short) use an interesting composite routing metric; see [Hed91] and [Far93] respectively as well as [Hui95]. In this section we discuss how the IGRP/EIGRP routing metric is composed and show that it is bounded but not monotonic.

IGRP was designed with a number of goals in mind. These included cycle free routing, fast response with low overhead, multipath routing and the ability to provide multiple types of service. In IGRP instead of a simple single metric, a set of path functions is maintained and this set is used to produced a composite metric. The composite metric is based upon four path functions and five constants. The four path functions are as follows:

1. Topological Delay (Distance)
2. Bandwidth (Flow)
3. Load
4. Reliability

Topological delay is the same as a distance metric. It is the sum of the transmission delays along the path to the root and represents the amount of time it would take a fixed size packet to reach the root assuming an unloaded network. Bandwidth is the minimum bandwidth encountered along the path to the root. It is expressed as an inverse. The delay and bandwidth for each edge are constants that depend upon the transmission medium. Load is the percentage of the available capacity that is already utilized. Reliability corresponds to the probability that a packet will arrive at the root. In addition to the above functions, a separate hop count is maintained.

The details of how load and reliability are computed are not provided in [Hed91]. One way to calculate the load correctly is to compute an effective bandwidth at each individual edge and then take the minimum over the edges along the path to the root.

The complete formula for the metric, as given in Cisco's documentation, is as follows:

$$(K1 * b + (K2 * b)/(256 - l) + K3 * d) * (K5/(r + K4))$$

where

$K1$, $K2$, $K3$, $K4$ and $K5$ are constants
b = inverse bandwidth
l = load: $0 \ldots 255$ (255 is saturated)
d = delay
r = reliability: $0 \ldots 255$ (255 is 100% reliable)

The path having the smallest composite metric is considered the best path. Consider that if delay decreases then the metric is reduced and likewise if bandwidth increases then the metric is reduced also. Furthermore if the load decreases or the reliability increases then the metric is reduced. Cisco's documentation does not indicate what are the permissible values for the constants. The default settings for $K1$, $K2$, $K3$, $K4$ and $K5$ are 1, 0, 1, 0, and 1 respectively. This yields the default formula:

$$b + d$$

While IGRP used a number of heuristics to prevent cycle formation, it could not guarantee tree maintenance. Enhanced IGRP is based on the same metric as IGRP, but replaces its heuristics with coordinated updates via diffusing computations. This protocol is documented in a paper available from Cisco [Far93] (see also [AGB94]), and its tree maintenance is based on the diffusing update algorithm (DUAL) of Garcia-Luna-Aceves [Gar89, Gar93]. The diffusing update algorithm was developed for shortest paths and the use of a nonmonotonic composite metric results in a different behaviour than what is described in the respective papers. Whereas shortest paths require a single diffusing computation, under the nonmonotonic composite metric of latency and bandwidth it is possible for multiple diffusing computations to occur in response to a single change in topology. This means that the protocol may have a higher overhead than is expected. Although it is counterintuitive, a decrease in the metric for one node can lead to an increase in the metric for another node [Sch97].

We conclude this section by analyzing the the composite metric of IGRP and EIGRP. This metric is bounded but it is not monotonic. Boundedness is easy to see. Consider that for each of the path functions, its contribution can only increase or maintain the composite metric and since we are minimizing along nondecreasing paths, the metric is bounded.

We will now show that the IGRP metric is not monotonic. Let us assume the default formula of $b + d$ which is the sum of the inverse bandwidth and the latency.

IGRP Metric

The composite *IGRP metric* $(W, wf, M, mr, met, \prec)$ is captured by the following:

1. $W \subseteq \{\langle b, d\rangle | b, d \in \{n | n \geq 0\}\}$
2. wf assigns to each edge an ordered pair from W
3. $M \subseteq \{\langle b, d\rangle | b, d \in \{n | n \geq 0\}\}$
4. $mr = \langle 0, 0\rangle$
5. $met(\langle b_1, d_1\rangle, \langle b_2, d_2\rangle) = \langle max(b_1, b_2), d_1 + d_2\rangle$
6. $\langle b_1, d_1\rangle \prec \langle b_2, d_2\rangle \Leftrightarrow (b_1 + d_1) < (b_2 + d_2)$

For clarity we have defined $\prec$ so that we are minimizing in the above definition. Consider $\langle b_1, d_1\rangle, \langle b_2, d_2\rangle \in M$ and $\langle b_3, d_3\rangle \in W$ such that $d_3 < d_2 < d_1 < b_1 < b_2 = b_3$. Assume that $\langle b_1, d_1\rangle \prec \langle b_2, d_2\rangle$. Applying *met* we get $met(\langle b_1, d_1\rangle, \langle b_3, d_3\rangle) = \langle b_3, d_1 + d_3\rangle$ and $met(\langle b_2, d_2\rangle, \langle b_3, d_3\rangle) = \langle b_3, d_2 + d_3\rangle$. However by $d_2 < d_1$, we get $(b_3 + d_1 + d_3) > (b_3 + d_2 + d_3)$ and thus $\langle b_3, d_1 + d_3\rangle \succ \langle b_3, d_2 + d_3\rangle$. Monotonicity does not hold in this case.

A more careful analysis is needed to show that IGRP is nonmonotonic in practice. A complete discussion of how the scaling is done for inverse bandwidth and delay is beyond the scope of this paper, but for most media the inverse bandwidth is the dominant factor of the two by orders of magnitude. Thus the cumulative addition of latency along a path is still not significant in comparison to the bandwidth of the path. Thus a node will minimize inverse bandwidth before latency and the above example will hold in practice.

6. Concluding Remarks

We developed a theory for deciding, for any routing metric and any network, whether the messages in this network can be routed along paths whose metric values are maximum. In order for the messages in a network to be routed along paths whose metric values are maximum, the network needs to have a rooted spanning tree that is maximal with respect to the routing metric. We identified two important properties of routing metrics: boundedness and monotonicity, and showed that these two properties are both necessary and sufficient to ensure that any network has a maximal tree with respect to any (bounded and monotonic) metric.

In related work [Sch97] we have shown that the distance-vector paradigm may be extended to arbitrary bounded and monotonic metrics such that a maximal tree will always be built. Furthermore, the presented protocol will still build a tree for any bounded and nonmonotonic metric such as the one used in IGRP.

We discussed how to combine two (or more) routing metrics into a single composite metric such that if the original metrics are bounded and monotonic, then the composite metric is bounded and monotonic. We then showed that the composite routing metrics used in IGRP (Inter-Gateway Routing Protocol) and EIGRP (Enhanced IGRP) are bounded but not monotonic.

Further investigation into the composition of maximizable routing metrics is a promising direction for future research.

References

[AB97] Y. Afek and A. Bremler, "Self-Stabilizing Unidirectional Network Algorithms by Power Supply," *Proceedings of the Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*,1997.

[AGB94] R. Albrightson, J.J. Garcia-Luna-Aceves, and J. Boyle, "EIGRP - a fast routing protocol based on distance vectors," *Proceedings of Networld/Interop 94*, May 1994.

[AG94] A. Arora and M. G. Gouda, "Distributed Reset," *IEEE Transactions on Computers*, Vol. 43, No. 9, September 1994, pp 1026-1039.

[BF89] B. Rajagopalan and M Faiman, "A New Responsive Distributed Shortest-Path Routing Algorithm," *ACM SIGCOMM 1989 Symposium*, September 1989.

[CD94] Zeev Collin and Shlomi Dolev, "Self-Stabilizing Depth First Search," *Information Processing Letters*, Vol. 49, pp. 297-301, 1994.

[CNRS97] E. Crawley, R. Nair, B. Balagopalan, and H. Sandick, "A Framework for QOS-based Routing in the Internet," *Internet Draft, draft-ietf-qosr-framework-02.txt*, November 1997.

[Far93] Dino Farinacci, "Introduction to Enhanced IGRP (EIGRP)," *http://www.cisco.com*, cisco Systems, July 1993.

[Gar89] J. J. Garcia-Luna-Aceves, "A Unified Approach to Loop-Free Routing Using Distance Vectors or Link States," *ACM SIGCOMM '89 Symposium*, September 1989.

[Gar93] J. J. Garcia-Luna-Aceves, "Loop-Free Routing Using Diffusing Computations," *IEEE/ACM Transactions on Networking*, Vol 1., No. 1, February 1993.

[GS94] Mohamed G. Gouda and Marco Schneider, "Stabilization of Maximum Flow Trees," Invited talk: *Proceedings of the Third Annual Joint Conference on Information Sciences*, November 1994.

[GS95] Mohamed G. Gouda and Marco Schneider, "Maximum Flow Routing," *Proc. of The Second Workshop on Self-Stabilizing Systems*, UNLV, pp. 2.1 - 2.13, 1995.

[Hed88] C. Hedrick, "Routing Information Protocol," *RFC 1058*, June 1988.

[Hed91] C. Hedrick, "An introduction to IGRP," Center for Computer and Information Services, Laboratory for Computer Science Research, Rutgers University, August 22, 1991. (Also at http://www.cisco.com).

[Hui95] C. Huitema, *Routing in The Internet*, Prentice Hall PTR, 1995.

[Sch93] M. Schneider, "Self-Stabilization," *ACM Computing Surveys*, Vol. 25, No. 1, March 1993.

[Sch97] M. Schneider, *Flow Routing in Computer Networks*, Ph.D. Dissertation, Department of Computer Sciences, The University of Texas at Austin, December 1997. URL: http://www.cs.utexas.edu/users/marco.

[WC95] Z. Wang and J. Crowcroft, "Bandwidth-Delay Based Routing Algorithms," *Proceedings of the 1995 IEEE Global Telecommunications Conference*, Singapore 1995.