

Stability of a Class of Dynamic Routing Protocols (IGRP) *

Steven Low

AT&T Bell Labs
600 Mountain Ave
Murray Hill, NJ 07974
slow@research.att.com

Pravin Varaiya

EECS Department
University of California
Berkeley, CA 94720
varaiya@helios.berkeley.edu

Abstract

We perform an exact analysis of the dynamic behavior of IGRP, an adaptive shortest-path routing algorithm widely used in the industry, on a simple ring network. The distance metric is a weighted sum of traffic-sensitive and traffic-insensitive delay components. We relate the optimality and stability of the protocol to the ratio of the weights. In particular, we show that if the traffic-insensitive component is not given enough weight, then starting from any initial routing, the subsequent routings after finitely many update periods will oscillate between two worst cases. Otherwise, the successive routings will converge to the unique equilibrium routing. We also show that load sharing among routes whose distances are within a threshold of the minimum distance help stabilize the dynamic behavior.

1 Introduction

A packet-switched network transports messages, packaged into streams of packets, between end-points. An end-point may represent a computer, a local area network, an audio or video source, or a database, etc. To each source-destination pair is usually associated a set of routes, or paths across the network. A routing strategy assigns to each packet a route from the set. Each route is assigned a cost called 'distance', which is usually a measure of hop-count or delay on the route. Under the shortest path routing, a packet is assigned the route with the least distance. Alternatively, load may be shared among multiple routes with similar distances. In static routing the decision is independent of

the traffic condition in the network. In dynamic routing the decision adapts to changing traffic condition and can potentially better balance the carried load. Typically the distance metric is defined to be a function of traffic on the route and route distances vary as the traffic condition fluctuates. Time is divided into update periods. At the end of each period, a new routing is computed using route distances in the current period, and then used in the next period to route packets for all source-destination pairs. The next routing thus depends on the current period's routing. Within a period packets between each source-destination pair follow a fixed route (except possibly for load sharing). Like any delayed feedback system it may be plagued by severe oscillation if not properly designed. Each node x stores a routing table, with an entry (z, w) for each destination node $z \neq x$. The entry (z, w) means that a packet arriving at node x that is destined for node z will be sent to x 's neighbor node w . The routing $y(n)$ in period n , $n = 0, 1, \dots$, may represent the routing tables at all nodes. We are interested in the 'stability' of $y(n)$ for a class of dynamic routing protocols.

The protocol we consider is modeled after IGRP [3] (Inter Gateway Routing Protocol), an adaptive routing algorithm developed by Cisco and used in many of its products. IGRP uses a variant of the updating mechanism in the standard Bellman-Ford algorithm [1, pp. 318-322]. The route distance that is exchanged periodically among neighboring nodes is changing rather than fixed. A simplified description is as follows. Periodically (every 90sec) each node x exchanges with its neighbors the distance $D(x, z)$ between itself and all other nodes $z \neq x$. $D(x, z)$ represents the shortest distance between node x and node z . Based on this information, each node x computes its shortest route to every other node z . It then broadcasts these new distances to its neighbors in the next update period, and the cycle repeats. Without load sharing the route between each source-destination pair

*Research supported by Pacific Gas and Electric Company. The authors are grateful to Shau-Ming Lun, Felix Wu, and Ning Xiao of U.C. Berkeley and Steve Callahan and Omid Razavi of PG&E.

is unique - one with the shortest distance. With load sharing the traffic between each source-destination pair may be distributed among more than one route.

The distance metric ¹ used in IGRP takes the form [3]

$$D(x, z) = k_i d_i(x, z) + k_s d_s(x, z)$$

where $k_i, k_s \geq 0$ are user-settable protocol parameters. The 'topological delay' $d_i(x, z)$ is a traffic-insensitive delay component that generalizes the notion of hop-count. Each link (k, m) between nodes k and m is assumed to have a fixed transmission capacity $\mu_{(k,m)} \geq 0$ in bits per second; $\mu_{(k,m)} = 0$ if there is no direct link between the nodes. A route $p = \{(l_1, l_2), (l_2, l_3), \dots, (l_{n-1}, l_n)\}$ is a sequence of links. The 'topological delay' $d_i(x, z)$ along route p connecting nodes x and z is $\sum_{l \in p} \mu_l^{-1}$, the total transmission time end-to-end for one bit of data. It is directly proportional to hop-count when each link has the same transmission capacity. Note, however, that the protocol is designed for networks in which propagation delay is negligible compared to transmission delay. $d_s(x, z)$ is a traffic-sensitive delay component which measures spare capacity of the route under the current routing (see (1-2) in §2).

Assuming that the traffic is stationary, we are interested in the dynamic behavior of IGRP. Specifically, we investigate the effect of the protocol parameters k_s, k_i on the 'stability' of $y(n)$, the routing in period n , and the optimality of the equilibrium routing when $y(n)$ is 'stable'. Our analysis provides insight in setting the protocol parameters. This problem was considered in [2]. The protocol there, however, is modeled after one used in the ARPANET and has a different distance metric from IGRP. We will compare our analysis and that in [2] after we have introduced our protocol model in §2. Following [2] we restrict attention to a simple ring network. We believe the intuition obtained from this analysis applies to more general topologies.

The paper is organized as follows. A model of IGRP is given in §2. §3 establishes the existence and optimality of the unique equilibrium routing. §4 investigates the dynamic behavior of the algorithm without load sharing. It is proved under certain conditions that if the traffic-insensitive component is not given enough weight, i.e. k_s/k_i too large, then starting from any initial routing, the subsequent routings after finitely many update periods will oscillate between two worst cases. On the other hand, if the traffic-insensitive

¹We have ignored a factor in the distance metric that measures the 'reliability' of a route.




Figure 1: A ring network

component is given sufficient weight, i.e. k_s/k_i sufficiently small, then regardless of the initial routing, the successive routings will converge to the unique equilibrium routing. IGRP also allows load sharing among routes whose distances are within a threshold of the shortest distance. We model this feature in §5 and show that load sharing has a stabilizing effect on the dynamic behavior of IGRP. Proofs of all our results can be found in [5] and [4].

2 Protocol Model

The network consists of an undirected ring. For analytical simplicity we consider a continuum of nodes on the ring, represented by $t \in [0, 1]$; see Figure 1. We assume node 1 (or equivalently node 0) is the only destination and every node $t \in (0, 1)$ has a source rate of $r(t)$ in bits per second. A routing $y \in [0, 1]$ takes the following simple form: under routing y , a node $t < y$ routes its traffic in the negative, or clockwise, direction and a node $t \geq y$ routes its traffic in the positive, or counter-clockwise, direction. Hence the routing decision at each node is simply to decide whether to send its traffic to the left or to the right neighbor. We assume time is slotted into update periods and all nodes operate synchronously. At the end of each update period, the nodes exchange update information, and compute their new routes for the next period.

The link capacity at node t in each direction is $\mu(t)$.

The topological delay at t is

$$d_i^-(t) = \int_0^t \mu^{-1}(s) ds$$

in the negative direction (to node 0), and

$$d_i^+(t) = \int_t^1 \mu^{-1}(s) ds$$

in the positive direction (to node 1). At the end of an update period with routing y , each node t computes the flow at t in the negative and the positive directions

$$f^-(t, y) = \int_t^y r(s) ds \mathbf{1}[t < y]$$

$$f^+(t, y) = \int_y^t r(s) ds \mathbf{1}[t > y]$$

The traffic-sensitive delay components, given by the reciprocal of the spare capacity of the route at t , are

$$d_s^-(t, y) = \frac{1}{\min_{s \in (0, t)} [\mu(s) - f^-(s, y)]} \quad (1)$$

$$d_s^+(t, y) = \frac{1}{\min_{s \in (t, 1)} [\mu(s) - f^+(s, y)]} \quad (2)$$

in the two directions. The shortest distance from t to the destination in the negative and the positive directions are respectively

$$D^-(t, y) = k_s d_s^-(t, y) + k_i d_i^-(t) \quad (3)$$

$$D^+(t, y) = k_s d_s^+(t, y) + k_i d_i^+(t) \quad (4)$$

for some $k_s, k_i \geq 0$. The new routing \hat{y} is the solution to

$$D^-(\hat{y}, y) = D^+(\hat{y}, y). \quad (5)$$

[2] considered the same ring network and analyzed the stability of an adaptive routing protocol used in the ARPANET. The metric $D^-(t, y)$ ($D^+(t, y)$) used there was the sum, or integral, over all nodes t of a given function $d(f^-(t, y))$ ($d(f^+(t, y))$) of flow $f^-(t, y)$ ($f^+(t, y)$) at t in the negative (positive) direction. It was proved there that if $d(0) = 0$, then the routing oscillates between two worst routes [2, Proposition 1], and that if $d(0)$ exceeds certain threshold, then the routing is stable [2, Propositions 5 and 6]. The metric used in IGRP is the sum of traffic-sensitive and traffic-insensitive components, weighted by the user-settable protocol parameters k_s, k_i . We show that the stability and optimality of the protocol are related to the ratio k_s/k_i . The same tradeoff between stability and responsiveness to traffic conditions manifest itself here as it did in [2].

3 Optimality of Equilibrium Routing

For the rest of this paper, we make the simplifying assumption that the transmission capacity is equal at all nodes, i.e. $\mu(t) \equiv 1$. We further assume for stability that $\int_0^1 r(s) ds < 1$. The first assumption reduces equations (1-4) to

$$D^-(t, y) = k_s d_s^-(y) + k_i d_i^-(t) \quad (6)$$

$$D^+(t, y) = k_s d_s^+(y) + k_i d_i^+(t). \quad (7)$$

where

$$d_s^-(y) = \frac{1}{1 - R(y)}$$

$$d_s^+(y) = \frac{1}{1 - R(1) + R(y)}$$

$$R(y) := \int_0^y r(s) ds$$

and

$$d_i^-(t) = t$$

$$d_i^+(t) = 1 - t.$$

Given the current routing y , the new routing \hat{y} is the solution to (5). A routing y^* is an equilibrium routing if

$$D^-(y^*, y^*) = D^+(y^*, y^*).$$

Proposition 1 *There exists a unique equilibrium routing.*

We next consider the optimality of the equilibrium routing. We will use the expression for the delay (sojourn time) through an $M/M/1$ queue with arrival rate R and service rate μ given by

$$\frac{1}{\mu - R}$$

to interpret the protocol. Under routing y , the delay at node $t < y$ is $1/[1 - f^-(t, y)]$ in the negative direction. This delay is maximum at the bottleneck node 0+ and is equal to

$$\tilde{W}^- = \frac{1}{1 - R(y)}$$

Similarly, the delay in the positive direction is maximum at the bottleneck node 1- and is equal to

$$\tilde{W}^+ = \frac{1}{1 - (R(1) + R(y))}$$

5c.4.3

for nodes $t \in [y, 1]$. Suppose, however, that each node $t \in (0, y)$ knows only its own preferred route (i.e. negative direction) and hence the preferred routes of nodes $s < t$, but does not know the value of current routing y and hence does not know the preferred route of nodes $s > t$. Then it would anticipate the delay at the bottleneck node $0+$ to be at least

$$W^-(t) = \frac{1}{1 - R(t)}, \quad t \in (0, y)$$

Similarly, a node $t \in [y, 1]$, knowing only its own preferred route but not the value of y , would anticipate the delay at the bottleneck node $1-$ to be at least

$$W^+(t) = \frac{1}{1 - (R(1) - R(t))}, \quad t \in [y, 1].$$

Hence the total delay at the bottlenecks is at least

$$\begin{aligned} W(y) &= \int_0^y W^-(t)dt + \int_y^1 W^+(t)dt \\ &= \int_0^y \frac{dt}{1 - R(t)} + \int_y^1 \frac{dt}{1 - R(1) + R(t)} \end{aligned}$$

It is reasonable to choose a routing y to minimize $W(y)$. Since $y = \frac{1}{2}$ is a good routing when traffic is uniform on the ring or when topological distance is the only cost ($k_s = 0$), we consider the following more general optimization:

$$\min_{y \in [0, 1]} W(y) + a(y - \frac{1}{2})^2 \quad (8)$$

where $a \geq 0$ is a given constant. a measures the relative weight we place on the two cost components $W(y)$ and $(y - 1/2)^2$. According to the following proposition it is related to the relative weight of the two delay components in (6) and (7): a larger a in the objective function (8) corresponds to a heavier weight on the traffic-insensitive components in the distance metric.

Theorem 2 *Let $y^*(k_s, k_i)$ denote the unique equilibrium routing with protocol parameters k_s, k_i . Then, $y^*(k_s, k_i)$ is the unique minimizer for (8) for every k_s, k_i satisfying*

$$\frac{k_s}{k_i} = a^{-1}$$

Under routing y the delay at the bottleneck $0+$ is in fact \tilde{W}^- for nodes $t \in (0, y)$, and the delay at the bottleneck $1-$ is in fact \tilde{W}^+ for nodes $t \in [y, 1]$. Hence a better optimization is

$$\min_{y \in [0, 1]} \tilde{W}(y) + a(y - \frac{1}{2})^2. \quad (9)$$

where

$$\begin{aligned} \tilde{W}(y) &= \int_0^y \tilde{W}^- dt + \int_y^1 \tilde{W}^+ dt \\ &= \frac{y}{1 - R(y)} + \frac{1 - y}{1 - R(1) + R(y)} \end{aligned}$$

It can be shown that solutions of (8) and of (9) coincide if and only if the minimizer y^* for (8) also satisfies

$$\frac{y^*}{1 - y^*} = [b(y^*)]^2$$

where

$$b(y) := \frac{1 - R(y)}{1 - R(1) + R(y)}$$

is the ratio of the spare capacity at the bottleneck nodes $0+$ and $1-$.

4 Stability Without Load Sharing

Recall that a new route is selected every update period by solving (5).

Proposition 3 *Given any routing $y \in [0, 1]$, a new routing $\hat{y} \in [0, 1]$ given by the solution of (5) exists and is unique if and only if*

$$\frac{k_s}{k_i} \leq \frac{1 - R(1)}{R(1)} \quad (10)$$

If the condition in Proposition 3 is not satisfied, i.e.

$$\frac{k_s}{k_i} > \frac{1 - R(1)}{R(1)}$$

then (5) has no solution \hat{y} in $[0, 1]$ if either $D^-(0, y) > D^+(0, y)$ or $D^-(1, y) < D^+(1, y)$. We naturally extend the new routing to be $\hat{y} := 0$ in the former case and $\hat{y} := 1$ in the latter case. Then starting with any routing y such that $D^-(0, y) > D^+(0, y)$ or $D^-(1, y) < D^+(1, y)$, the subsequent routings will oscillate between 0 and 1, the worst possible scenario. In fact we can give a more complete characterization of the dynamic behavior of the protocol.

For $y \in [0, 1]$, let

$$\begin{aligned} f(y) &:= \frac{k_s}{k_i} (d_s^-(y) - d_s^+(y)) \\ &= \frac{k_s}{k_i} \cdot \frac{2R(y) - R(1)}{(1 - R(y))(1 - R(1) + R(y))} \\ g(y) &:= d_i^+(y) - d_i^-(y) = 1 - 2y \end{aligned}$$




Figure 2: Construction of y_i and z_i

Then (5) is equivalent to

$$f(y) = g(\hat{y}) \quad (11)$$

and condition (10) is equivalent to $f(1) \leq g(0) = 1$. When $f(1) > 1$ there are $y \in [0, 1]$ such that $f(y) > 1 = g(0)$ or $f(y) < -1 = g(1)$ for which no solution \hat{y} to (11) exists in $[0, 1]$. We define the new routing to be $\hat{y} := 0$ in the former case and $\hat{y} := 1$ in the latter case. The unique equilibrium routing y^* in Proposition 1 satisfies $f(y^*) = g(y^*)$.

Suppose $f(1) > g(0)$. Define the sequences $(y_n, n \geq 1)$ and $(z_n, n \geq 0)$ by

$$\begin{aligned} z_0 &= 1 \\ y_{n+1} &= \inf\{y \mid f(y) = z_n\}, \quad n \geq 0 \\ z_{n+1} &= g(y_{n+1}), \quad n \geq 0 \end{aligned} \quad (12)$$

The construction is illustrated in Figure 2. f is continuous and increasing; it is invertible if it is strictly increasing. In that case

$$y_{n+1} = f^{-1}(z_n)$$

When convenient, we may abuse notation and denote (12) by $y_{n+1} = f^{-1}(z_n)$. Let y^* and z^* be such that

$$f(y^*) = g(y^*) = z^*$$

Lemma 4 Suppose $\frac{k_s}{k_i} > \frac{1-R(1)}{R(1)}$. Then

$$\begin{aligned} y_1 &> y_3 > \dots \geq y^* \geq \dots > y_4 > y_2 \\ z_0 &> z_2 > \dots \geq z^* \geq \dots > z_3 > z_1 \end{aligned}$$

Since both $(y_{2n-1}, n \geq 1)$ and $(y_{2n}, n \geq 1)$ are monotone and bounded, the limits

$$\bar{y} := \lim_n y_{2n-1} \quad \text{and} \quad \underline{y} := \lim_n y_{2n}$$

exist and $\underline{y} \leq y^* \leq \bar{y}$. Moreover,

$$\begin{aligned} f(y_{2n-1}) &= z_{2n-2} = g(y_{2n-2}) \\ f(y_{2n}) &= z_{2n-1} = g(y_{2n-1}) \end{aligned}$$

and hence the continuity of f and g implies that

$$\begin{aligned} f(\bar{y}) &= g(\bar{y}) \\ f(\underline{y}) &= g(\underline{y}) \end{aligned}$$

This means that starting with initial routing $y_1 = \bar{y}$ or $y_1 = \underline{y}$, the subsequent routings will oscillate between \bar{y} and \underline{y} .

Theorem 5 Suppose $\frac{k_s}{k_i} > \frac{1-R(1)}{R(1)}$. Let y_1 be an initial routing.

1. If $y_1 = \underline{y}$ or $y_1 = \bar{y}$, then subsequent routings oscillate between \underline{y} and \bar{y} .
2. If $y_1 < \underline{y}$ or $y_1 > \bar{y}$, then subsequent routings after finitely many update periods oscillate between 0 and 1.
3. If $\underline{y} < y_1 < \bar{y}$, then subsequent routings converge to the unique equilibrium routing y^* , provided

$$\frac{k_s}{k_i} < \frac{(1-R(1))^2}{r(\bar{t})} \quad (13)$$

where $\bar{t} := \arg \max_{t \in (\underline{y}, \bar{y})} r(t)$.

If we strengthen condition (13) by changing \bar{t} to

$$\bar{t} := \arg \max_{t \in [0, 1]} r(t)$$

then we obtain the following corollary, which corresponds to the case in which $\underline{y} = 0$ and $\bar{y} = 1$.

Corollary 6 Given any initial routing $y_1 \in [0, 1]$, the successive routings converge to the unique equilibrium routing y^* , provided

$$\frac{k_s}{k_i} < \frac{(1-R(1))^2}{r(\bar{t})}$$

where $\bar{t} := \arg \max_{t \in (0, 1)} r(t)$.

On the other extreme, we have the following situation, which corresponds to the case in which $\underline{y} = y^* = \bar{y}$.

5c.4.5

Corollary 7 Given any initial routing $y_1 \in [0, 1]$ with $y_1 \neq y^*$, the subsequent routings after finitely many update periods oscillate between 0 and 1, provided

$$\frac{k_s}{k_i} > \frac{1}{r(\underline{t})}$$

where $\underline{t} := \arg \min_{t \in (0,1)} r(t)$.

Recall that the route distance in IGRP has a traffic-sensitive and a traffic-insensitive components. In our case, the traffic-insensitive component measures the distance between the source and the destination, and the traffic-sensitive component measures the delay at the bottleneck nodes under the current routing. The route distance is the sum of these two components, weighted by k_s and k_i . Corollary 7 says that if the traffic-insensitive component is not given enough weight, then starting from any initial routing, the subsequent routings after finitely many update periods will oscillate between two worst cases. If the static component is given a large enough weight, then according to Corollary 6, starting from any initial routing, the successive routings will converge to the unique equilibrium routing. Such stability is achieved, however, at the cost of reduced adaptivity to traffic conditions.

5 Stability With Load Sharing

IGRP extends the basic algorithm modeled in the previous section to allow load sharing among several routes. To route traffic from node x to node z , node x computes the distance to z via each neighbor. Instead of routing all its traffic to z via the neighbor on a shortest route, node x splits the traffic among routes whose distances are within a threshold of the minimum distance. We show in this section that load sharing has a stabilizing effect on the dynamic behavior of IGRP since it enlarges the stability region on the protocol parameters.

We model load sharing in our ring network as follows: under routing $y \in [0, 1]$, a node $t < y - \epsilon$ routes its traffic in the negative direction, a node $t > y + \epsilon$ routes its traffic in the positive direction, and a node $y - \epsilon \leq t \leq y + \epsilon$ splits its traffic equally in both directions. See Figure 3. Then under routing y , the flow at t in the negative direction is modified to $f^-(t, y) = \int_t^{y-\epsilon} r(s)ds + \int_{y-\epsilon}^{y+\epsilon} \frac{1}{2}r(s)ds$ for $t < y - \epsilon$, and $f^-(t, y) = \int_t^{y+\epsilon} \frac{1}{2}r(s)ds$ for $y - \epsilon \leq t < y + \epsilon$. In the positive direction, the flow becomes $f^+(t, y) = \int_{y+\epsilon}^t r(s)ds + \int_{y-\epsilon}^{y+\epsilon} \frac{1}{2}r(s)ds$ for $t > y + \epsilon$,




Figure 3: Network model with load sharing

and $f^+(t, y) = \int_{y-\epsilon}^t \frac{1}{2}r(s)ds$ for $y - \epsilon < t \leq y + \epsilon$. Hence, the traffic-sensitive delays become

$$d_s^-(y) = \frac{1}{1 - \frac{1}{2}(R(y - \epsilon) + R(y + \epsilon))}$$

$$d_s^+(y) = \frac{1}{1 - R(1) + \frac{1}{2}(R(y - \epsilon) + R(y + \epsilon))}$$

with the understanding that $R(-\epsilon) = R(0) = 0$ and $R(1 + \epsilon) = R(1)$. As before, given the current routing y , the new routing \hat{y} is the solution to (5). With load sharing, conclusions similar to those in the previous section can be drawn with less stringent stability conditions. They are summarized in the following propositions, which include the previous ones as special cases ($\epsilon = 0$).

Proposition 8 There exists a unique equilibrium routing under load sharing with any parameter $\epsilon \geq 0$.

Denote $\epsilon_0 := \frac{1}{2}R(\epsilon)$ and $\epsilon_1 := \frac{1}{2}[R(1) - R(1 - \epsilon)] = \int_{1-\epsilon}^1 \frac{1}{2}r(s)ds$.

Proposition 9 Given any routing $y \in (0, 1)$, the new routing \hat{y} given by (5) exists and is unique if and only if k_s/k_i is less or equal to the minimum of

$$\frac{(1 - \epsilon_0)(1 - R(1) + \epsilon_0)}{R(1) - 2\epsilon_0} \quad \text{and} \quad \frac{(1 - \epsilon_1)(1 - R(1) + \epsilon_1)}{R(1) - 2\epsilon_1}$$

It can be verified that the condition in Proposition 3 implies that in the above proposition.

Proposition 10 Given any initial routing y_1 , the successive routings converge to the equilibrium routing y^* , provided

$$\frac{k_s}{k_i} < \frac{2}{r(\bar{t}) \left[\frac{1}{(1-R(\bar{t})+\epsilon_0)^2} + \frac{1}{(1-R(\bar{t})+\epsilon_1)^2} \right]} \quad (14)$$

where $\bar{t} := \arg \max_{t \in (0,1)} r(t)$.

The stability condition in Corollary 6 implies the condition

$$\frac{k_s}{k_i} < \frac{(1-R(1) + (\epsilon_0 \wedge \epsilon_1))^2}{r(\bar{t})} \quad (15)$$

which in turn implies (14) in Proposition 10, where $(\epsilon_0 \wedge \epsilon_1)$ denotes $\min\{\epsilon_0, \epsilon_1\}$. Hence load sharing enlarges the stability region on the protocol parameters.

Proposition 11 Given any initial routing $y_1 \in (0, 1)$ with $y_1 \neq y^*$, the subsequent routings after finitely many update periods oscillate between 0 and 1, provided

$$\frac{k_s}{k_i} > \frac{1}{r(\underline{t})} \cdot \frac{2}{\frac{1}{(1-\epsilon_0)^2} + \frac{1}{(1-\epsilon_1)^2}}$$

where $\underline{t} := \arg \min_{t \in (0,1)} r(t)$.

6 Conclusion

IGRP is widely used in practice. We have performed an exact analysis of its dynamic behavior on a simple ring network. It provides insight in setting the protocol parameters k_s, k_i . The distance metric in IGRP is the sum of traffic-sensitive and traffic-insensitive delay components, weighted by k_s, k_i . We have related the optimality (Theorem 2) and the stability (Corollaries 6 and 7) of the protocol to the ratio of these parameters. Roughly, the routing will converge to the unique equilibrium routing if the traffic-insensitive component is given sufficient weight; otherwise, it will oscillate between two worst cases after finitely many update periods.

References

- [1] D. Bertsekas and R. Gallager. *Data Networks*. Prentice-Hall Inc., 1987.
- [2] Dimitri P. Bertsekas. Dynamic behavior of shortest path routing algorithms for communication networks. *IEEE Transactions on Automatic Control*, pages 60–74, February 1982.

- [3] Charles L. Hedrick. An introduction to IGRP. preprint, Laboratory for Computer Science Research, Rutgers University, New Jersey, August 1991.
- [4] S. Low and P. Varaiya. Dynamic behavior of a class of adaptive routing protocols (IGRP). in preparation, 1993.
- [5] Steven Low. *Traffic Management of ATM Networks: Service Provisioning, Routing, and Traffic Shaping*. PhD thesis, UC Berkeley, May 1992.