

Routing Using Potentials: A Dynamic Traffic-Aware Routing Algorithm

Anindya Basu
Bell Laboratories

basu@research.bell-labs.com

Alvin Lin
MIT

alvinl@mit.edu

Sharad Ramanathan
Bell Laboratories

sharadr@physics.bell-labs.com

Abstract

We present a routing paradigm called PB-routing that utilizes steepest gradient search methods to route data packets. More specifically, the PB-routing paradigm assigns scalar potentials to network elements and forwards packets in the direction of maximum positive force. We show that the family of PB-routing schemes are loop free and that the standard shortest path routing algorithms are a special case of the PB-routing paradigm. We then show how to design a potential function that accounts for traffic conditions at a node. The resulting routing algorithm routes around congested areas while preserving the key desirable properties of IP routing mechanisms including hop-by-hop routing, local route computations and statistical multiplexing. Our simulations using the *ns* simulator indicate that the traffic aware routing algorithm shows significant improvements in end-to-end delay and jitter when compared to standard shortest path routing algorithms. The simulations also indicate that our algorithm does not incur too much control overheads and is fairly stable even when traffic conditions are dynamic.

Categories & Subject Descriptors: C.2.2 Routing Protocols.

General Terms: Algorithms.

Keywords: Congestion, Potential, Routing, Steepest Gradient, Traffic Aware.

1. Introduction

Routing mechanisms in the Internet have typically been based on shortest-path routing for best effort traffic. This often causes traffic congestion, especially if bottleneck links on the shortest path severely restrict the effective bandwidth between the source and the destination.

Traditionally, congestion control in the Internet has been provided by end-to-end mechanisms. An example is the TCP congestion control mechanism that works by adjusting the sending rate at the source when it detects congestion at a bottleneck link (for details, see [25]). If multiple traffic streams share the same bottleneck link, each gets only a fraction of the bottleneck link bandwidth even though there may be bandwidth available along alternate paths in the network. Moreover, queueing delays at the bottleneck link can add significantly to end-to-end delays. Finally, varying traffic conditions can make this queueing delay variable, thereby adding to jitter.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'03, August 25–29, 2003, Karlsruhe, Germany.

Copyright 2003 ACM 1-58113-735-4/03/0008 ...\$5.00.

One way to address the problem of end-to-end delay and jitter is to use traffic engineering (TE) techniques in conjunction with circuit-based routing. In this case, the routing process assumes that the traffic demands between source-destination pairs are known a priori and computes end-to-end paths (circuits) satisfying the traffic demands. End-to-end circuits are then set up along the computed paths using a resource reservation protocol such as RSVP [13]. Data packets are now source routed along these pre-computed paths.

There are some drawbacks to this kind of circuit-based routing. First, if traffic sources are bursty (which is mostly the case in the Internet — see, for example, [21]), resources may be reserved unnecessarily, thereby negating the benefits of statistical multiplexing. Second, traffic demands between network nodes are hard to estimate a priori. Also, if traffic patterns change, it is possible that a global recomputation is necessary to determine the most optimal routing. Third, when future demands are unpredictable, it is difficult to route current demands such that a future demand has the maximum chance of being routed successfully without requiring the rerouting of existing demands. Indeed, it can be shown that this problem is NP-hard even for the simplest cases [12].

In this paper, we present an alternate methodology for traffic-aware routing that is based on steepest gradient search methods. We call this methodology *potential based routing*, or *PB-routing*. It preserves the hop-by-hop routing philosophy of the Internet, and does not require a priori knowledge of traffic demands between network nodes. At the same time, PB-routing is able to route packets around the congested hot-spots in the network by utilizing alternate routes that may be non-optimal. This reduces end-to-end delays and jitter and increases the bandwidth utilization in the network. Since packets are not source routed, PB-routing can adapt to changes in traffic conditions without requiring any global recomputation of routes. Furthermore, end-to-end resource reservation is not required — hence, the benefits of statistical multiplexing are still available. Note that PB-routing can only provide performance improvements of a statistical nature, and not explicit worst-case bounds on delay and jitter (as opposed to TE techniques). However, our simulations indicate that there is significant improvement in delay and jitter for most traffic streams when PB-routing is used.

The key idea in PB-routing is to define a scalar field on the network, which is used to define a potential on every network element (NE). The routing algorithm at each NE now computes the route to the destination as the direction (i.e., the next hop) in which the potential field decreases fastest (direction of maximum force or steepest gradient). We show that by assigning the NE potentials differently, a whole family of routing algorithms can be designed. For example, the standard shortest path algorithm can be shown to be a special case of PB-routing if the potential at each NE is set to be a linear, monotonically increasing function of the shortest distance from the NE to the destination. The routing algorithm can be made traffic-aware by setting the poten-

tial at each NE to be a weighted sum of the shortest path potential and a metric that represents the traffic potential at the NE. We show how to define such a metric later in the paper.

More intuitively, the PB-routing algorithm views the entire network as a terrain, with many negotiable obstacles created by congestion. The high obstacles represent areas of the network with high congestion (and therefore, high potential). The idea is to find a path from the source to the destination by avoiding the high obstacles as much as possible.

Since PB-routing depends on traffic information at the various NEs, it is necessary that this information be disseminated efficiently without compromising end-to-end packet delays and jitter. In our simulations, we adapt a link-state routing protocol for this purpose. We also describe an optimization that significantly reduces the control overheads of this protocol without sacrificing the performance of the routing algorithm.

The main contributions of this paper are as follows. While steepest gradient search methods have been well-studied, the novel idea in this paper is the design of a potential field for traffic-aware routing that guarantees desirable properties such as loop-free routing. We demonstrate that our design of a potential function, which is a hybrid of traffic metrics and link costs, ensures that packets avoid congested areas but do not traverse the network using random walks. In fact, how far the path of a packet deviates from the standard shortest path can be controlled by a configurable parameter. In our simulations, we have observed significant improvements in end-to-end delay and jitter over a variety of networks and traffic conditions without requiring too much control overheads. We believe that the general framework of PB-routing could be adapted for optimizing various other metrics through careful design of potential functions. This is especially true of overlay networks (see, for example, [2]) where application specific metrics that require optimization could be converted into appropriate potential functions.

The rest of the paper is organized as follows. In the next section, we describe the PB-routing model in greater detail and prove some of its properties analytically. This is followed by Sections 3 and 4 where we describe our implementation of traffic-aware routing using the PB-routing paradigm and evaluate its performance. We then describe related work in Section 5 and conclude in Section 6.

2. The PB-routing Paradigm

In this section, we present the key theoretical ideas underlying the PB-routing paradigm. We emphasize here that the PB-routing paradigm represents a *family of network routing algorithms*. Hence, we first provide a description of the generic PB-routing algorithm. Using this generic formulation, we prove the key properties that are common to all the algorithms in the PB-routing family. We then describe two specific instantiations of the PB-routing paradigm. The first instantiation is the standard shortest path algorithm — we refer to this as **SPP**. We then show how **SPP** can be modified to be traffic-aware (called the **PBTA** algorithm) and analyze its desirable properties.

System Model. In order to describe the PB-routing paradigm, we first need to define some terminology and a system model. We model a network of nodes connected by bidirectional links as a directed graph $G = (N, E)$. The set of nodes in the network is represented by the set of vertices N in G . Similarly, the set of edges E in G corresponds to the set of links in the network, where e_{uv} is a directed edge from vertex u to vertex v with cost metric c_{uv} that is strictly positive. Since the network links represented by the edges in E are bidirectional, it is easy to see that if edge $e_{uv} \in E$, then $e_{vu} \in E$. For the rest of this paper, we shall use the terms nodes (links) and vertices (edges) interchangeably. Each node v can act as a traffic source and/or sink.

Furthermore, every node v has a set of $Z(v)$ neighbors denoted by $nbr(v)$. Thus, the indegree and outdegree of any node v are both equal to $Z(v)$.

2.1 Routing with Potentials

The PB-routing paradigm defines a scalar field on the network over which packets are routed. The potential at any node v is a function of v and the destination d for which we need to find a route. More formally, with each node v (and destination d), we associate a potential $V^d(v)$ that is *single-valued*. Note that if the destination d changes, the potential function for v changes as well. We prove all the properties of PB-routing assuming that the destination d is fixed. Since the potential functions for different destinations are independently defined, it follows that our assumption about a fixed destination is not restrictive. For the rest of this paper, we shall use $V(v)$ to denote the potential at a node v when the destination is clear from the context.

Now consider a packet p at a node v whose destination is node d . In order to reach d , p must be forwarded to one of the $Z(v)$ neighbors of v . To determine this “next hop” neighbor, we define a “force” on the packet p at v based on the potentials at v and its neighbors. For a neighbor $w \in nbr(v)$, we can define the force $F_{v \rightarrow w}$ as the discrete derivative of V with respect to the link metric as

$$F_{v \rightarrow w} = \frac{(V(v) - V(w))}{c_{vw}} \quad (1)$$

The packet p is now directed to the neighbor $x \in nbr(v)$ for which the force $F_{v \rightarrow x}$ is maximum and positive. In other words, each packet follows the direction of the steepest gradient downhill to reach its destination. We now prove the following general property of the PB-routing paradigm.

THEOREM 2.1. *The PB-routing paradigm is loop-free if the potential function $V(v)$ is time invariant.*

Proof: We prove this by contradiction. Consider a packet p that is routed along a closed loop on the network, beginning and ending at node v . Let this closed loop be the directed path $v = v_0 \rightarrow v_1 \rightarrow v_2 \rightarrow \dots \rightarrow v_{k-1} \rightarrow v_0 = v$. For p to be routed along this path, the work done defined by the forces in equation (1) must be strictly positive. This is because the routing algorithm always directs packets in the direction of the maximum positive force. More formally,

$$\sum_{i=0}^{k-1} F_{v_i \rightarrow v_{(i+1) \bmod k}} \cdot c_{v_i v_{(i+1) \bmod k}} > 0 \quad (2)$$

Using equation (1), we get

$$\sum_{i=0}^{k-1} V(v_i) - V(v_{(i+1) \bmod k}) > 0 \quad (3)$$

Since $V(v)$ is a time invariant, single valued function of v , the LHS of equation (3) must be identically zero, which is a contradiction. Hence, the PB-routing paradigm is loop-free as long as the potential function is single-valued and time invariant. ■

Now consider any packet p at a node v . Since p always moves in the direction of the maximum positive force, p will be forwarded to a neighbor w that satisfies the following

$$V(v) - V(w) > 0 \quad (4)$$

However, this is not possible if v is a local minima,¹ i.e., we have

$$\forall w \in nbr(v) \ V(v) - V(w) < 0 \quad (5)$$

¹The term *local minima* means that the potential of p is lower than that of any of its neighbors.

In other words, p will get stuck at a node v if v is not the destination but is a local minima. This implies the following

LEMMA 2.2. *If the potential function has a minimum only at the destination and no other local minima, all packets will eventually reach their destinations.*

2.2 The SPP Algorithm using PB-routing

We now describe **SPP** in detail. In the traditional shortest path algorithms, a packet traverses the shortest path from source s to destination d . Typically, shortest paths can be computed using well known algorithms (e.g. Dijkstra's shortest path algorithm [7]).

Now, let p be a packet at a node v going to destination d , and let D_{vw} be the length of the shortest path in the network graph connecting nodes v and w . We set the potential at node v to be

$$V^d(v) = V(D_{vd}) \quad (6)$$

where $V(x) = ax + b, a > 0$, is a single-valued, monotonically increasing, linear function of x . To reach destination d , a packet p at node v selects the next hop $w \in nbr(v)$ such that the force

$$F_{v \rightarrow w} = \frac{V(D_{vd}) - V(D_{wd})}{c_{vw}} \quad (7)$$

is maximum and positive. It is now easy to see that **SPP** is loop-free by Theorem 2.1 in the absence of topology changes since the potential function V is single-valued and time invariant everywhere.

To show that each packet p eventually gets to its destination if **SPP** is used, we first prove the following lemma.

LEMMA 2.3. *The potential function V has no local minima.*

Proof: The proof is by contradiction. Let v be a node that has a local minima. In other words, we have

$$\forall w \in nbr(v) \ V(D_{vd}) < V(D_{wd}) \quad (8)$$

Since v is a monotonic increasing function of D_{wd} , equation (8) implies that

$$\forall w \in nbr(v) \ D_{vd} < D_{wd} \quad (9)$$

Now let u be the next hop on the shortest path from v to d . Then, using the properties of the shortest path computation algorithms, we have

$$D_{vd} = c_{vu} + D_{ud} \quad (10)$$

where c_{vu} represents the cost metric for the link e_{vu} . However, since $u \in nbr(v)$, using equation (9), we have

$$D_{vd} < D_{ud} \quad (11)$$

Using equations (10) and (11), we conclude that $c_{vu} < 0$ which contradicts the assumption that link metrics are strictly positive. ■

Now consider the destination d . Using the fact that link metrics are strictly positive, and that V is a monotonically increasing function, we have $V(D_{dd}) = V(0) < V(D_{vd})$, where $v \in N$ and $v \neq d$. We therefore conclude that the potential function has a minimum at the destination, and no other local minima. Using Lemma 2.2, we assert that every packet p is guaranteed to eventually reach its destination.

Finally, we show that **SPP** does indeed route using the shortest path. To prove this property, we use the following lemma.

LEMMA 2.4. *For any node v , and destination d , if the next hop computed by the shortest path algorithm is u , then the next hop computed by **SPP** is also u .*

Proof: Wlog, let $w \in nbr(v)$ be such that $w \neq u$. Then, by the shortest path property, we have

$$c_{vu} + D_{ud} = D_{vd} \leq c_{vw} + D_{wd} \quad (12)$$

This implies that

$$\frac{D_{vd} - D_{wd}}{c_{vw}} \leq 1 \quad (13)$$

and

$$\frac{D_{vd} - D_{ud}}{c_{vu}} = 1 \quad (14)$$

Using equations (13) and (14), we get

$$\frac{D_{vd} - D_{ud}}{c_{vu}} \geq \frac{D_{vd} - D_{wd}}{c_{vw}} \quad (15)$$

Using equation (15) and the fact that $V(x)$ is a monotonically increasing linear function of x , we conclude that

$$\frac{V(D_{vd}) - V(D_{ud})}{c_{vu}} \geq \frac{V(D_{vd}) - V(D_{wd})}{c_{vw}} \quad (16)$$

In other words, the force in the direction of u is maximum and positive. Therefore, **SPP** chooses u as the next hop. Note that we make the implicit assumption here that if there are multiple paths with the same minimum cost, both algorithms use the same deterministic procedure to break ties. ■

COROLLARY 2.5. *Let V be of the form $V(x) = ax + b, a > 0$. Then, for any node v , and a node $u \in nbr(v)$, we have $\frac{V(D_{vd}) - V(D_{ud})}{c_{vu}} \leq a$.*

It is now easy to see that **SPP** simulates the standard shortest path routing. We know that both the algorithms compute the same next hop at every node for every packet with destination d . Thus, every packet from source s to destination d follows the same sequence of links in both cases. Therefore, we have the following property

THEOREM 2.6. ***SPP** correctly simulates the standard shortest path routing algorithms.*

2.3 Generalizing Potentials to be Traffic-Aware

We now show how the PB-routing paradigm can be used to construct traffic-aware routing algorithms. In order to do this, we have to design a potential that includes a traffic component. For the purposes of this paper, we use the outgoing queue sizes at a network node v as a measure of traffic at that node.² In the rest of this section, we describe the design of the traffic potential in greater detail followed by proofs of its relevant properties.

2.3.1 Design of the traffic potential.

In order to design a traffic potential, we first introduce some more notation. Let Q_{vu} denote the queue length on the outgoing link e_{vu} adjacent to node v in the original network graph. The quantity Q_{uv} is defined similarly. Let BW_{vu} be the bandwidth associated with e_{vu} — we assume that $BW_{vu} = BW_{uv}$ and $c_{vu} = c_{uv}$, where c_{vu} is the cost metric associated with the link e_{vu} in the network graph. Finally, let the normalized queue length on a link e_{vu} in the original graph be

$$q_{vu} = \frac{Q_{vu}}{BW_{vu}} \quad (17)$$

The normalized queue length on a link represents the time it will take for the current queue on that link to drain. For the rest of this paper, we use the terms *queue length* and *normalized queue length* interchangeably.

²Note that instead of the outgoing queue sizes, some other metric may also be used as a measure of traffic.




Figure 1: (a) A network represented by a directed graph, and (b) the corresponding transformed graph over which the traffic potentials are defined. The square boxes in (b) are the nodes that correspond to the edges in (a). The lines with double arrows represent two directed edges.

Graph Transformation. In order to design a traffic potential based on queue lengths, we need to take into account the fact that queues exist on network links and not on network nodes. This is because the queues for buffering packets reside on the linecards in most modern routers. For simplicity, we assume that there is a 1-1 mapping between router linecards and links. For this purpose, we define a transformation on the original graph $G = (N, E)$ that represents the actual network. Let the transformed graph be $G' = (N', E')$, where N' and E' are defined as follows

$$N' = \{e_{vu} | e_{vu} \in E\} \quad (18)$$

$$E' = \{(e_{vu}, e_{uw}) | (e_{vu} \in E \wedge e_{uw} \in E)\} \quad (19)$$

Thus, the nodes of the transformed graph G' are the directed edges of the network graph G . The edges of the transformed graph represent common nodes between edges in the original graph. Two edges in the original graph have a common node iff the head of the first edge is the tail of the second edge. It is easy to see that any node e_{vu} in N' has an outdegree $Z(u)$, and an indegree $Z(v)$. The generalized degree $Z(e)$ of the node $e = e_{vu}$ is defined as

$$Z(e) = Z(e_{vu}) = \frac{Z(u)}{c_{vu}} \quad (20)$$

We denote the maximum generalized degree of a node in G' as Z_{max} .

A network graph and the corresponding transformed graph are shown in Figures 1(a) and (b), respectively. Each square box in the transformed graph corresponds to an edge in the original graph. The edges in the transformed graph obey the rules in (19). The circles in Figure 1(b) are the nodes from the original graph that have been superposed — they are not part of the transformed graph. The dotted lines connect a superposed node to each of its outgoing edge nodes in the transformed graph.

Next, we define a matrix operator A on the transformed graph G' as follows

$$A_{vu, xw} = \begin{cases} \frac{1}{c_{vu}} & \text{if } u = x \\ 0 & \text{if } u \neq x \end{cases} \quad (21)$$

The matrix element between the node e_{vu} and the node e_{uw} is the inverse of the cost metric for the edge e_{vu} in the original network graph.³ Note that this matrix A is not symmetric.

We are now ready to define a traffic potential Φ_{vu} on each node e_{vu} of the transformed graph G' . This means that the traffic potential is

³If there are multiple edges connecting v to u , there will be separate rows in A for each such edge.

really defined on every edge e_{vu} of the original network graph G . We require that the traffic potential at each node e_{vu} in G' be the maximum of q_{vu} and the solution to the discrete laplace's equation (see [17])

$$\sum_{w \in nbr(u)} A_{vu, uw} (\Phi_{uw} - \Phi_{vu}) + (Z_{max} - Z(e_{vu}))(0 - \Phi_{vu}) = 0 \quad (22)$$

The physical interpretation of the second term in the sum is that for every vertex e_{vu} on the new graph G' with generalized degree $Z(e_{vu}) \leq Z_{max}$, there are $Z_{max} - Z(e_{vu})$ ghost nodes connected to e_{vu} with edges of unit cost. We require that the value of the scalar field at these ghost nodes is zero, which defines the boundary conditions for the discrete laplace's equation above. The equation (22) then has a unique solution (see [14] for a proof).

We now present a more intuitive picture of the traffic potential. The traffic potential function corresponds to the surface of a taut elastic membrane that covers the network like a tent. The queues on the links can be thought of as vertical poles that hold up the membrane, which is “pegged down” at the ghost nodes. This implies that the tent surface is “propped up” by the larger queues and the smaller queues do not touch the tent surface. Hence the potential at a node e_{vu} in the transformed graph is the larger of the solution to equation (22) and the normalized queue length on the edge e_{vu} in the original network graph. More formally,

$$\Phi_{vu} = \max \left(\frac{1}{Z_{max}} \sum_{w \in nbr(u)} A_{vu, uw} \Phi_{uw}, q_{vu} \right) \quad (23)$$

Let $r(v)$ be the ratio of the maximum link cost metric to the minimum link cost metric among all the outgoing links adjacent to node v . Then, we can define the potential Φ_v at a node v in the original network graph as

$$\Phi_v = \max \left(\frac{1}{Z(v)} \sum_{w \in nbr(v)} \Phi_{vw}, \frac{r(v)\Phi_{max} + \Phi_{min}}{r(v) + 1} \right) \quad (24)$$

where Φ_{max} is the maximum traffic potential on an outgoing link adjacent to node v , and Φ_{min} is the minimum. In other words, the traffic potential on any node in the original network graph G is the maximum of two quantities — the average of the potentials on the outgoing edges of the node, and a weighted average of the maximum and minimum potentials on the outgoing edges of the node.

2.4 PBTA — Traffic-Aware Routing with Potentials

We now describe how to route packets on the network graph using the traffic potential. Consider a packet p at node v with destination d . We define an effective potential on the graph that combines the effect of traffic load with the standard shortest path routing algorithm. The value of this potential at node v given by

$$\mathcal{V}(v) = (1 - \alpha)V(D_{vd}) + \alpha\Phi_v \quad (25)$$

and its value on the edge e_{vu} given by

$$\mathcal{V}(vu) = (1 - \alpha)V(D_{ud}) + \alpha\Phi_{vu} \quad (26)$$

where V is the shortest distance potential function defined in equation (6), and Φ_{vu} and Φ_v and defined by the equations (23) and (24). The parameter α ($0 < \alpha < 1$) sets the relative weights of the traffic potential Φ and the shortest distance potential V . We can then define a “force” on the packet p at node v towards a neighbor $u \in nbr(v)$ as

$$F_{v \rightarrow u} = \frac{\mathcal{V}(v) - \mathcal{V}(vu)}{c_{vu}} \quad (27)$$

This equation can be rewritten as

$$F_{v \rightarrow u} = (1 - \alpha)F_{spp}(v, u) + \alpha F_t(v, u) \quad (28)$$

where $F_{spp}(v, u) = \frac{V(D_{vd}) - V(D_{ud})}{c_{vu}}$ is the shortest path component, and $F_t(v, u) = \frac{\Phi_v - \Phi_{vu}}{c_{vu}}$ is the traffic component. The packet at v is then transmitted on the edge e_{vw} towards which the total force is maximum and positive.

2.5 Properties of the routing Algorithm

If the traffic patterns (and hence, the queue sizes) are stationary, and the network topology does not change, then the single valued potential function \mathcal{V} is time invariant. Thus, by Theorem 2.1, the **PBTA** algorithm is loop free. We address the issue of dynamic traffic patterns in Section 2.6.

We now show that the **PBTA** algorithm is stable. Following [26], we define stability to mean that no single queue length grows unbounded independently. The stability property ensures that the network load is evenly distributed — no single queue is allowed to become a bottleneck. To prove this property, we first prove the following lemma.

LEMMA 2.7. *Let the potential function for shortest path be of the form $V(x) = ax + b$, $a > 0$. Then, $|F_{spp}(v, u)| \leq a$ for any pair of nodes v, u such that $u \in nbr(v)$.*

Proof: From the definition of $F_{spp}(v, u)$, we know that

$$F_{spp}(v, u) = \frac{V(D_{vd}) - V(D_{ud})}{c_{vu}} \quad (29)$$

From Corollary 2.5, we know that $F_{spp}(v, u) \leq a$. Furthermore, we know by the shortest path property (applied to node u) that

$$D_{ud} \leq c_{uv} + D_{vd} \quad (30)$$

Using the fact that $c_{uv} = c_{vu}$, we have

$$-1 \leq \frac{D_{vd} - D_{ud}}{c_{vu}} \quad (31)$$

This, together with equation (29), and the fact that $V(x) = ax + b$ implies that

$$-a \leq \frac{V(D_{vd}) - V(D_{ud})}{c_{vu}} = F_{spp}(v, u) \quad (32)$$

Therefore $|F_{spp}(v, u)| \leq a$. ■

THEOREM 2.8. *The queue length on any single link never grows without bound independently.*

Proof: We prove this by contradiction. Consider a link e_{vu} that has the largest queue length, say q_{max} . Let q_{max} be much higher compared to the queue length q_{wx} on any other link e_{wx} such that the resulting traffic potential at all links except e_{vu} satisfy the condition $\Phi_{wx} > q_{wx}$. Thus q_{max} is large enough such that no other queue besides itself has any effect on the traffic potential.

To add to the queue length q_{max} on link e_{vu} , packets must be sent along the link e_{vu} by node v to some destination d (say). Now, we know that

$$F_{v \rightarrow u} = (1 - \alpha)F_{spp}(v, u) + \alpha F_t(v, u) \quad (33)$$

as defined earlier. We observe that $F_t(v, u)$ is really the slope of the surface corresponding to the traffic potential field. The force $F_t(v, u)$ from v to u must be negative since the largest queue lies in the direction from v to u . Hence, the maximum value of $F_t(v, u)$ is given by

$$F_t(v, u) \leq -\frac{q_{max}}{D_{max}} \quad (34)$$

where D_{max} is the maximum shortest path distance between any two nodes in the network. Combining equations (33) and (34), we have

$$F_{v \rightarrow u} \leq (1 - \alpha)F_{spp}(v, u) - \alpha \frac{q_{max}}{D_{max}} \quad (35)$$

Using Lemma 2.7, we have

$$F_{v \rightarrow u} \leq (1 - \alpha)a - \alpha \frac{q_{max}}{D_{max}} \quad (36)$$

Thus, if $q_{max} > \left(\frac{1-\alpha}{\alpha}\right) a D_{max}$, then the net force from v to u is negative. This means that every packet is directed away from the queue from v to u when the queue grows sufficiently large. Therefore, no queue on the network can grow unbounded independently. ■

Finally, we show that the effective potential field has no minima. This implies that no packet ever gets stuck at any node except the destination.

THEOREM 2.9. *There is no node v in a graph (except the ghost nodes) where the force on a packet due to the combined shortest path and traffic potentials is negative in all directions.*

Proof: Consider a node v in the network graph that is not a ghost node. Let the shortest path potential function V be of the form $V(x) = ax + b$, $a > 0$. By definition, and using equation (27), the force from v in the direction of e_{vu} is

$$F_{v \rightarrow u} = (1 - \alpha)F_{spp}(v, u) + \alpha F_t(v, u) \quad (37)$$

Now consider the shortest path direction. It is possible to show that if u is the next hop node in that direction, we have

$$F_{v \rightarrow u} = (1 - \alpha)a + \alpha F_t(v, u) \quad (38)$$

If this force is positive, v is not a local minima. Otherwise, we have

$$(1 - \alpha)a + \alpha \frac{\Phi_v - \Phi_{vu}}{c_{vu}} < 0 \quad (39)$$

Since $(1 - \alpha)a$ is positive, the second term must be negative, and we have

$$(1 - \alpha)a < \alpha \frac{\Phi_{vu} - \Phi_v}{c_{vu}} \quad (40)$$

Now consider the link e_{vw} such that

$$\Phi_{vw} = \min_{x \in nbr(v)} \Phi_{vx} = \Phi_{min} \quad (41)$$

Then the total force in the direction of e_{vw} is

$$F_{v \rightarrow w} = (1 - \alpha)F_{spp}(v, w) + \alpha \frac{\Phi_v - \Phi_{min}}{c_{vw}} \quad (42)$$

Using Lemma 2.7, we can show that

$$F_{v \rightarrow w} \geq -(1 - \alpha)a + \alpha \frac{\Phi_v - \Phi_{min}}{c_{vw}} \quad (43)$$

Using equation (40), the definition of $r(v)$, and the definition of Φ_v , we have

$$\begin{aligned} & -(1 - \alpha)a + \alpha \frac{\Phi_v - \Phi_{min}}{c_{vw}} \\ & > \alpha \left(\frac{\Phi_v - \Phi_{min}}{c_{vw}} - \frac{\Phi_{vu} - \Phi_v}{c_{vu}} \right) \\ & > \frac{\alpha}{c_{vw}} ((\Phi_v - \Phi_{min}) - r(v)(\Phi_{vu} - \Phi_v)) \\ & \geq 0 \end{aligned}$$

Therefore, the force $F_{v \rightarrow w}$ is positive, and v is not a local minima, which proves the theorem. ■

2.6 Stability Issues

Earlier, we have proved that the **PBTA** algorithm is stable (implicitly) assuming that the queue length information is instantaneously available at all nodes whenever a change occurs. In other words, the stability properties that we have proved may not hold if network nodes have stale (queue length) information. Since the packet propagation delay across a network is finite, it is possible that very rapid changes in queue lengths will cause instabilities in the routing algorithm. However, in our simulations (described in Section 4), we have observed that the **PBTA** algorithm continues to be stable. We now provide some physical arguments as to why this is the case.

The nature of the **PBTA** algorithm is such that the effect of a large queue on the traffic potential is spread out over the neighbors of the queue. Consider, for instance, a graph with nodes on a regular square lattice, with exactly one queue of size q . Then, the scalar field Φ at a point at distance d from the queue satisfies $\Phi \sim q \ln(\frac{1}{d})$. Similarly, for a regular cubic lattice, we have $\Phi \sim \frac{q}{d+1}$. Thus, at least for regular graphs, we see that the scalar potential Φ due to a queue decays slowly with distance from the queue. While it is difficult to compute such expressions for arbitrary graphs, we postulate that the nature of the decay is qualitatively the same. A slow rate of spatial decay means that a large queue will influence the shape of the potential surface even at points that are not in the immediate neighborhood.

Following the tent analogy, the large queues are the high “poles” that poke into the taut elastic tent fabric, propping it up and determining the shape of the surface. In contrast, smaller queues (i.e. queues for which $q_{vu} < \frac{1}{z_{max}} \sum_{w \in nbr(v)} A_{vw,wx} \Phi_{wx}$) do not touch the tent fabric and therefore have no effect on the contour of the potential surface.

We also observe that the relative changes in large queues occur at a slow rate — more specifically, for a large queue of size q , $\frac{\delta q}{q}$ changes slowly since q is large. Thus, we can say that *the shape of the potential surface is almost wholly determined by the large queues, which change relatively slowly.* For the most part, we have observed that the rate of this change is slower than the mean packet (carrying queue length information) propagation delay across the network.

Furthermore, the *flooding optimization* described in Section 3 shows that the network can (in effect) be partitioned into “sub-domains”. A key property of such a partition is that the queue lengths on links outside the sub-domain do not affect the traffic potential field inside the sub-domain. Given that a sub-domain is typically smaller than the whole network, the mean packet propagation delay across a sub-domain is much smaller than the time required to flood queue length information across the whole network. Hence, it is even more unlikely that the shape of the potential surface changes faster than it takes for a packet (containing queue length information) to traverse a sub-domain (on an average). In other words, it is very improbable that the traffic potential field computed using (possibly stale) queue length information at a node would be significantly different from the field computed with the latest queue length information. Therefore, it is highly unlikely that the **PBTA** algorithm will exhibit instability.

In summary, we note that a key requirement for the **PBTA** algorithm to be stable is that *the traffic potential surface changes slower than the time it takes for a packet to traverse a sub-domain*. While this could be viewed as a limitation, we have informally argued above that for most realistic network scenarios, such is not the case. This is reflected in our simulations, which show stable performance even in the face of stale queue length information. Of course, a more formal analysis of the regimes where the **PBTA** algorithm is stable for arbitrary network topologies and traffic matrices is intractable.

2.7 Other Traffic-Aware Routing Algorithms

It is possible to define other traffic-aware routing algorithms as alternatives to the **PBTA** algorithm. However, we are not aware of any

such algorithm that provides the following key benefits:

- Provable stability properties — no single queue height diverges independently when the **PBTA** algorithm is used. As we have argued in Section 2.6, the stability properties hold in a dynamic setting so long as the mean packet propagation delay within a “domain” is smaller than the rate at which the potential surface within a domain changes.
- We provide a *physical* criterion for deciding which queue lengths are relevant for the route computation process, without using any heuristic rules. Consequently, we are able to develop a flooding optimization (see Section 3) that limits the regions over which queue length information has to be propagated. This reduces overheads significantly without compromising performance.

We provide a more concrete illustration of the advantages of the **PBTA** algorithm by comparing it to a heuristic that uses shortest path routing where the link metrics are set to be proportional to the queue lengths on the links. A detailed analytical examination of the properties of this heuristic is beyond the scope of this paper. However, qualitatively speaking, the queue length based approach has the following disadvantages:

Bottlenecks. Since the link metrics in shortest path computations are additive, a path with many medium queues may be rejected in favor of a path with mostly small queues and a single large one. This means that individual links may become bottlenecks causing the corresponding queue lengths to diverge independently. In contrast the **PBTA** algorithm is able to distribute the network load more evenly over all queues.

Avoiding Congested Areas. The queue length algorithm determines the link metrics based solely on the queue lengths on the link, and does not take into account neighboring links. Consequently, packets may avoid congested links by minimal deviations around them. Therefore, the packets do not avoid areas of congestion from far enough away (as in the **PBTA** algorithm, where the effect of large queues can be felt far enough away). As a result, large queues on congested links can take longer to drain, causing bottlenecks to last longer.

2.8 Limitations

We now discuss some limitations of the **PBTA** algorithm. The first limitation is related to how fast the senders transmit packets. Clearly, if the sending rates are low compared to the link capacities, queues do not build up significantly. Therefore, no significant improvements in end-to-end delays is observed. On the other extreme, if sending rates are extremely high compared to link capacities, it is possible that all links in the network get saturated, and there are no alternate paths left for packets to traverse. In such a case also, the **PBTA** algorithm fails to improve performance. While it is difficult to estimate (for arbitrary network topologies) the range within which the improvements shown by the **PBTA** algorithm are significant, we have actually been able to verify this using simulations.

The second limitation relates to the nature of end-to-end delays. We know that the **PBTA** algorithm improves performance by attempting to route around large queues in the network. Therefore, significant improvements in end-to-end delays will not be observed in networks where queueing delays are very small compared to link latencies, such as in satellite networks. Furthermore, if the link propagation delays are very long, the routing algorithm may not be able to adapt to changing network conditions fast enough, thereby causing instabilities (see Section 2.6 for details).

Finally, if the network graph is sparsely connected (a linear graph in the worst case), the **PBTA** algorithm does not perform any better than

the standard shortest path algorithm. This is because typically, there are no alternate paths to the destination. In some cases it is possible that there is a bottleneck link which *must* be traversed in order to get to the destination, i.e., a cut-edge connecting two biconnected components. In such a case, if the link has a very high queue size, there are two possible alternatives for a packet wishing to traverse the link. The first possibility is to add to the queue irrespective of its size, and potentially get dropped. The other possibility is to go into a “holding pattern” (by traveling in some other direction away from the destination) till the bottleneck queue drains. In practice, the choice of the alternative depends on the setting of α , the maximum queue size, and the TTL (time to live) parameter of a packet. These can be tuned to achieve the desired behavior.

3. Implementation

In this section we describe our implementation of the **PBTA** algorithm using the *ns* simulator and elaborate on some of the implementation issues that arose. The *ns*-based implementation is used to evaluate the performance of the **PBTA** algorithm in the next section.

The implementation of the **PBTA** algorithm is based on a link-state routing protocol (such as OSPF [19]). Routers running a link-state protocol compute routes to different destinations based on information obtained from a link-state database. Each router (i.e., network node) has a copy of the link-state database that reflects the state of the network — the copies of the link-state database at the different routers are kept consistent by the routing protocol. Whenever the network state changes (for example, if a link fails), the network node that detects the change floods this information across the entire network by encapsulating it in a link state advertisement packet (or LSA). When new information is received, a network node recomputes routes to various destinations (typically) using a shortest path algorithm.

Our implementation modifies both the route computation algorithm and the information dissemination process. The route computation algorithm uses the link metric information and the queue length information in the link state database to compute routes, as described in Section 2.4. The information dissemination process propagates queue length information in addition to other link-state information. This raises three main issues: when and how frequently are link state updates (especially, queue length updates) sent out, how is the parameter α (the parameter that assigns the relative weights of the shortest path and traffic potentials) determined for a network, and how far is a given link state update (carrying queue length information) propagated.

Frequency of Update Information. The link-state update process is “trigger-based”. Whenever a link/node failure or recovery occurs, a new LSA is flooded across the network. In addition, each node monitors the queue length associated with each of its outgoing links. An LSA is sent out whenever the relative change in queue length exceeds a pre-configured threshold value, q_f , where $0 < q_f < 1$.

In addition, when a link state update is scheduled for transmission, the packet encapsulating it is always placed at the head of the appropriate outgoing queue. Thus, control packets are given priority over data packets in the implementation. This ensures that LSAs are disseminated across the network in a timely manner without being slowed down by heavily congested queues. Such a practice is fairly common in modern routers.

Configuring α . The setting of α is more involved. Obviously, α depends to some extent on the network topology. We now describe a criterion that can be applied without detailed knowledge of the network topology.

Let the shortest path distance between maximally separated nodes in the graph be D_{max} . In response to a change δq at the queue q on




Figure 2: Choice of α : (a) shows a potential surface when α is low – the SPP metric has more significance here which is reflected by the higher tilt of the plane. (b) shows a potential surface when α is high — the tilt is lower since the traffic potential has more significance. The dark line on each surface shows a route avoiding the high peaks, which are areas of high congestion.

link e , the force due to the traffic potential at a neighboring node u will change at least by

$$\delta F \sim \alpha \frac{\delta q}{D_{max}} \quad (44)$$

This worst-case scenario is achieved if the graph is effectively one-dimensional. For example, in an effectively two dimensional graph, we have a smaller change in F given by $\delta F \sim \alpha \frac{\delta q}{D_{max} \ln D_{max}}$.

Using equations (25) and (26), we can conclude that the maximum force is totally dominated by the traffic component (no matter what the shortest path related potential is) if

$$\delta F \sim \alpha \frac{\delta q}{D_{max}} \gg (1 - \alpha) \quad (45)$$

This means that the packets can get routed in a fashion that tries to avoid large queues as much as possible, without any regard to what the shortest path is. Hence, α should be set such that $\alpha \frac{\delta q}{D_{max}} \approx (1 - \alpha)$. We therefore choose α to be

$$\alpha = \alpha_0 \cdot \frac{D_{max}}{\delta q + D_{max}} \quad (46)$$

where $0 < \alpha_0 < 1$ is an initial value that is scaled by a factor given by the rest of the expression in (46). In our implementation, we set $\alpha_0 = 0.33$. The effect of choosing different values of α on the potential surface is shown in Figure 2. In our simulations we found that for a large variety of networks, α_0 can take on values over a fairly wide range (between 0 and 1) and still provide (close to) optimal performance. Therefore, it does not seem necessary to carefully tune the initial value α_0 by trial and error for a given network.

Distributing the Update Information. We now describe a way to reduce the control overheads by restricting how far a particular queue length update is sent from the originating node. We call this optimization the *flooding optimization*.

To understand qualitatively why such an optimization works, consider the behavior of the traffic potential Φ due to the queue on some link e . As explained earlier (see Section 2.6), this potential Φ decays (albeit slowly) with distance d from the actual queue. If indeed, as we postulate earlier, this decay obeys the power law, we can say that the change in traffic potential $\delta \Phi$ at a distance d from the queue whose length has changed by δq scales as $\frac{\delta q}{d^\gamma}$. Thus, the change in force δF , derived from the expression in equation (27) is given by

$$\delta F \sim (1 - \alpha)C + \alpha \frac{\delta q}{d^\gamma} \quad (47)$$

