Plasma Anypath Routing in Wireless Mesh Networks

Rafael Laufer, Luiz Filipe M. Vieira, Mario Gerla, and Leonard Kleinrock
Computer Science Department
Univeristy of California at Los Angeles

July 31, 2009
Technical Report UCLA-CSD-TR090022

***Abstract*—We present a new routing paradigm that generalizes opportunistic routing for the multi-gateway case. In plasma anypath routing, each packet is delivered over the best available path to one of the gateways. The choice of the path and gateway for each packet is not made beforehand by the source node, but rather on-the-fly by the mesh routers as the packet traverses the network. To our knowledge, the problem of gateway anycasting via anypath routing has not been explored before. We provide a theory capable of jointly optimizing the transmission rate and the set of next hops to reach the best subset of gateways. We propose an optimal distributed routing algorithm as well as a load-balancing technique to disperse the network traffic among multiple gateways. We validate our proposal with traces from an 802.11b testbed. Our results show that plasma anypath routing outperforms multirate anypath routing, with a maximum gain of 31% for two gateways and 64% for four gateways. We also show that the load can be distributed among the gateways and that plasma anypath routing is robust to wireless link fluctuations over long periods of time.**

I. Introduction

Routing in wireless multihop networks is still a challenge due to the high loss rate and dynamic quality of wireless links [1], [2]. Opportunistic or anypath routing has been recently proposed as a way to circumvent the wireless shortcomings by using multiple next hops for each destination [3]–[6]. In anypath routing, each packet is broadcast to a forwarding set composed of several neighbors, and the packet must be retransmitted only if none of the neighbors in the set receive it. Therefore, while the link to a given neighbor is down or performing poorly, another nearby neighbor may receive the packet and forward it on. This is in contrast to single-path routing where only one neighbor is assigned as the next hop for each destination. In this case, if the link to this neighbor is not performing well, a packet may be lost even though other neighbors may have overheard it. Anypath routing takes advantage of these opportunities to avoid unnecessary retransmissions, increasing the overall network performance.

Existing work on anypath routing [7] has focused only on a unicast delivery model, where each mesh node sends traffic to a single gateway. Usually, each node selects the closest gateway and sends all of its traffic to this particular gateway [8]. Albeit straightforward, this strategy can lead to both unfairness and under-utilization. In the case of a hotspot, a particular gateway may be overloaded with mesh nodes competing for a small bandwidth share. At the same time, another nearby gateway may be free and under-utilized.

In *plasma anypath routing*, we address this problem by allowing each packet to be delivered to any of the gateways. Packets also take one of the many available paths towards each gateway, so we take advantage of both anypath and anycast routing. Our idea is that a mesh node should not send its traffic to a single gateway. It is hard to accurately know the load and wireless conditions on the path to that gateway beforehand. A better approach is to let the network decide both the path and the gateway on-the-fly as the packet traverses the network. Ideally, a mesh node should just notify the network about its intention to send a packet to the Internet, and the network should be responsible for actually delivering it. Our main inspiration comes from a plasma lamp [9], where an inner electrode irradiates filaments towards the outer glass sphere. Those filaments are in fact electric currents flowing through high-conductivity regions. In our analogy, packets flow from a source node towards one of the gateways through low-noise and low-interference areas.

In this paper, we address the problem of jointly determining the forwarding set, transmission rate, and gateway subset for every node, such that the cost of every node to the Internet is minimized. We call this the *shortest plasma anypath problem*. To our knowledge, the problem of anycasting via anypath routing has not been considered before. We introduce a polynomial-time distributed routing algorithm to this problem and prove its optimality. The shortest plasma anypath problem is *no harder* than the well-known shortest-path problem, being therefore suitable for implementation at current routing protocols. We also introduce a load balancing scheme that network operators can use to easily shift the load from one gateway to the others.

We validate our proposal with traces from an 18-node 802.11b testbed of embedded Linux devices. Our results show that plasma anypath routing improves the end-to-end transmission time as we increase the number of deployed gateways. For two gateways, the expected trasnmission time is reduced by up to 31% and, with four gateways, the end-to-end transmission time can be up to 64% lower. We also show that it is straightforward to balance the load among the different gateways. Additionally, we show that plasma routing is very robust over time and not affected by channel variations. Therefore, having an up-to-date picture of the topology does not provide a significant benefit, which allows the routing protocol overhead to be reduced.

The remainder of the paper is organized as follows. In Section II, we review the theory of multirate anypath routing. In Section III, we introduce plasma anypath routing, present the proposed routing algorithm and prove its optimality. Results from the performance evaluation of the algorithm are presented in Section IV. Section V covers the related work and finally Section VI concludes the paper.

II. Multirate Anypath Routing

We introduced multirate anypath routing in [7]. In this section we review the theory of multirate anypath routing. Our main contributions are presented later in Section III.

A. Overview

In anypath routing, a node broadcasts a packet to *multiple* next hops simultaneously. Therefore, if the transmission to one neighbor fails, another neighbor which received the packet can forward it on. We define this set of multiple next hops as the *forwarding set* and we usually use J to represent it throughout the paper. A different forwarding set is used to reach each destination, in the same way a distinct next hop is used for each destination in classic routing. In *multirate* anypath routing, a node also uses a fixed bit rate r to transmit to each forwarding set. This allows nodes to take advantage of forwarding sets that can sustain a higher transmission rate as well as sets that can only communicate at a lower rate. For each destination, a node then maintains the forwarding set J and the transmission rate r that must be used to reach this set.

When a packet is broadcast to the forwarding set, there is a chance that more than one node receives the same packet. To avoid unnecessary duplicate forwarding, only one of these nodes should forward the packet on. For this purpose, each node in the set has a priority in relaying the received packet. A node only forwards a packet if all higher priority nodes in the set failed to do so. Higher priorities are assigned to nodes with lower costs to the destination. As a result, if the node with the lowest cost in the forwarding set successfully received the packet, it forwards the packet to the destination while others suppress their transmission. Otherwise, the node with the second lowest cost forwards the packet, and so on. A reliable anycast scheme [10] is necessary to enforce this relay priority and we talk more about this in Section II-B. The source rebroadcasts the packet until someone in the forwarding set receives and acknowledges it or a threshold is reached. Once a neighbor receives the packet, it repeats the same procedure until the packet is delivered to the destination.

Since we now use a set of next hops to forward packets, every two nodes are connected through a mesh composed of the union of multiple paths, with each node transmitting at a selected rate. Figure 1 depicts this scenario where nodes use a given bit rate to forward packets to a set of neighbors. The forwarding set is defined by the multiple bold arrows leaving each node. We define this union of paths between two nodes, with each node using a potentially different bit rate as a *multirate anypath*. In the figure, the anypath shown in bold is composed by the union of 11 different paths between the source s and destination d. At every hop, only a single node of the set forwards the packet on. Consequently, every packet from s traverses only one of the available paths to reach d. We show a path possibly taken by a packet using dashed lines. We use different dash lengths to represent the different transmission rates used by each node. A shorter dash represents a shorter time to send a packet, and thus a higher transmission rate. Succeeding packets may take completely different paths with other transmission rates along the way; hence the name multirate anypath. The path taken is determined on-the-fly, depending on which nodes of the forwarding set successfully receive the packet at each hop.

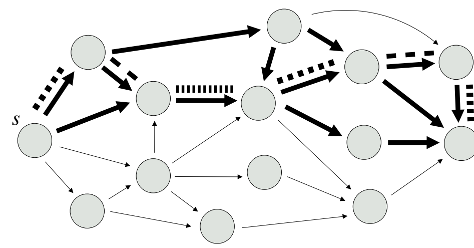

Figure 1. A multirate anypath connecting nodes s and d is shown in bold arrows. A packet from s traverses one of these paths to reach d, such as the path shown with dashed lines. Different dash lengths represent the different bit rates used by each node, with a shorter dash for higher rates.

B. System Models and Assumptions

We model the wireless mesh network as a hypergraph. A hypergraph $\mathcal{G} = (V, \mathcal{E})$ is composed of a set V of vertices or nodes and a set $\mathcal{E}$ of hyperedges or hyperlinks. A hyperlink is defined as an ordered pair (i, J), where $i \in V$ is a node and J is a nonempty subset of V composed of neighbors of i. Let R be the set of available bit rates for nodes to transmit. For each hyperlink $(i, J) \in \mathcal{E}$, we have a delivery ratio $p_{iJ}^{(r)}$ and a cost $d_{iJ}^{(r)}$ associated with each transmission rate $r \in R$. This reflects the fact that in wireless mesh networks, we have different delivery ratios and costs for each rate. If the set J has a single element j, then we just use j instead of J in our notation. In this case, $p_{ij}^{(r)}$ and $d_{ij}^{(r)}$ denote the link delivery ratio and cost at rate r, respectively.

The hyperlink delivery ratio $p_{iJ}^{(r)}$ is defined as the probability that a packet transmitted from i using rate $r \in R$ is successfully received by at least one of the nodes in J. One would expect that the receipt of a packet at each neighbor is correlated due to noise and interference. However, we conducted experiments which suggest that the loss of a packet at different receivers occur independently for light load regimes [7], which is also consistent with other studies [11]. With the assumption of independent losses, we have $p_{iJ}^{(r)} = 1 - \prod_{j \in J} \left(1 - p_{ij}^{(r)}\right)$.

Several MAC protocols have been proposed to guarantee the relay priority in the forwarding set [10]. Such protocols use different strategies for this purpose, such as time-slotted access, prioritized contention, and frame overhearing. Reliable anycast is an active area of research [10] and we assume that such a mechanism is in place to make sure that the relay priority is respected. The priority scheme parameters implemented in the MAC are derived from routing information. The details of the MAC, however, are abstracted from the routing layer. Practical routing protocols only incorporate the link delivery ratios into the routing metric in order to abstract from the MAC details [12], [13] and we take the same approach. The only MAC aspect that is important is the effectiveness of the relaying node selection. As long as the relaying node is actually the one with the lowest cost to the destination, there should be no significant impact on the routing performance.

C. Multirate Anypath Cost

We are interested in calculating the cost $D_i^{(r)}$ from a node i to a given destination via forwarding set J when i transmits at rate r. The cost $D_i^{(r)}$ is defined as the sum of two terms $D_i^{(r)} = d_{iJ}^{(r)} + D_J^{(r)}$, the hyperlink cost $d_{iJ}^{(r)}$ from i to J and the remaining cost $D_J^{(r)}$ from J to the destination. We now explain each one of these terms.

The metric used in multirate anypath routing is the expected anypath transmission time (EATT). When using EATT, the hyperlink cost $d_{iJ}^{(r)}$ for rate $r \in R$ is defined as

$$d_{iJ}^{(r)} = \frac{1}{p_{iJ}^{(r)}} \times \frac{s}{r}, \tag{1}$$

where $p_{iJ}^{(r)}$ is the hyperlink delivery ratio, s is the maximum packet size, and r is the bit rate. The hyperlink cost $d_{iJ}^{(r)}$ is basically the time it takes to transmit a packet of size s at a bit rate r over a lossy hyperlink with delivery ratio $p_{iJ}^{(r)}$. The EATT metric is a generalization of the expected transmission time (ETT) metric [13] used in single-path wireless routing.

The remaining cost $D_J^{(r)}$ is defined as a *weighted average* of the costs of the nodes in the forwarding set as

$$D_J^{(r)} = \sum_{j \in J} w_{ij}^{(r)} D_j, \text{ with } \sum_{j \in J} w_{ij}^{(r)} = 1, \tag{2}$$

where $D_j = \min_{r \in R} D_j^{(r)}$ and the weight $w_{ij}^{(r)}$ is the probability of node j being the relaying node. For example, let $J = \{1, 2, \ldots, n\}$ with costs $D_1 \leq D_2 \leq \ldots \leq D_n$. We refer to the probability $p_{ij}^{(r)}$ simply by p_j for convenience. Node j will be the relaying node only when it receives the packet and none of the nodes closer to the destination also receives it. This happens with probability $p_j(1-p_{j-1})(1-p_{j-2})\ldots(1-p_1)$. The weight $w_{ij}^{(r)}$ is then

$$w_{ij}^{(r)} = \frac{p_{ij}^{(r)} \prod_{k=1}^{j-1} \left(1 - p_{ik}^{(r)}\right)}{1 - \prod_{k \in J} \left(1 - p_{ik}^{(r)}\right)}, \tag{3}$$

with the denominator being the normalizing constant.

As an example of cost calculation, consider the network depicted in Figure 2. Figure 2(a) shows the link delivery ratios when the bit rate is 1 Mbps and Figure 2(b) shows the same ratios for 2 Mbps. For 1500-byte packets and a fixed transmission rate of 1 Mbps, the cost via J in Figure 2(a) is calculated as

$$\begin{aligned} D_i^{(1)} &= d_{iJ}^{(1)} + D_J^{(1)} \\ &= \frac{12,000 \text{ bits/1 Mbps}}{1-(1-0.33)(1-0.25)} + \frac{(0.25)36 + (0.75)(0.33)60}{1-(1-0.33)(1-0.25)} \\ &= 24 + 48 = 72 \text{ ms.} \end{aligned} \tag{4}$$

The dark gray arrows in the figure represent the 1-Mbps hyperlinks. For 2 Mbps, the lowest cost is via node j, as shown by the light gray arrows in Figure 2(b). The cost is calculated as $D_i^{(2)} = d_{ij}^{(2)} + D_j^{(2)} = 40 + 40 = 80$ ms. It is obvious from this example that increasing the rate does not always decrease the cost. The lowest cost at 2 Mbps is 80 ms, while at 1 Mbps it is 72 ms.

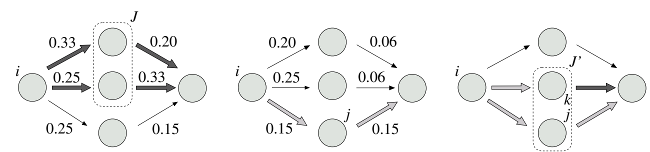

(a) 1-Mbps anypath (b) 2-Mbps anypath (c) Multirate anypath

Figure 2. A multirate anypath cost calculation example. The weights are the link delivery ratios at (a) 1 Mbps and (b) 2 Mbps. The multirate anypath in (c) yields the lowest cost by fixing the rate of nodes i and j at 2 Mbps (light gray arrows) and the rate of node k at 1 Mbps (dark gray arrow).

Instead of using a fixed rate for the entire network, a much better strategy is to take advantage of the good links of each rate. Considering the multirate approach in Figure 2(c), the lowest cost is via J' and the the best rate selection is for nodes i and j to transmit at 2 Mbps (light gray arrows) and for node k to transmit at 1 Mbps (dark gray arrow). The cost via the forwarding set J' is then calculated as

$$\begin{aligned} D_i &= d_{iJ'} + D_{J'} \\ &= \frac{12,000 \text{ bits/2 Mbps}}{1-(1-0.25)(1-0.15)} + \frac{(0.25)36 + (0.75)(0.15)40}{1-(1-0.25)(1-0.15)} \\ &= 16.6 + 37.2 = 53.8 \text{ ms.} \end{aligned} \tag{5}$$

Clearly, the multirate anypath offers a significantly lower cost to the destination. The estimated end-to-end transmission time with multirate is 53.8 ms, approximately 25% and 33% lower than the cost at 1 Mbps and 2 Mbps, respectively.

III. Plasma Anypath Routing

Although multirate anypath routing has many advantages, it is still limited by its unicast or point-to-point delivery model. With a unicast model, the usual approach for Internet access in wireless mesh networks is for each node to forward packets along the shortest path to the closest gateway [8]. However, there are a couple of issues with this method. First, it does not take advantage of the multiple paths available between the mesh node and the gateway to avoid noise and interference, as anypath routing does. Packets usually follow the same path to reach the gateway, and delivery can be severely affected if the quality of just one of the links drops. Additionally, using a single path also makes packets more susceptible to intra-flow interference, possibly resulting in lower throughput rates. Second, if the load is not evenly distributed across the network, we run into known fairness and under-utilization problems [14]. The worst-case scenario occurs when most of the nodes are grouped near one gateway (i.e., a hotspot). These nodes compete for a small fraction of the gateway's bandwidth, while nodes associated with other gateway may achieve a much larger throughput, leading to unfairness. If other gateways only have a few or no nodes associated with them, their bandwidth may also remain largely under-utilized.

Our key idea is that a mesh node should not be associated with a single gateway, but rather with multiple gateways in order to benefit from the full aggregate bandwidth, load balancing, and fairness. Additionally, nodes should not try to calculate in advance how much traffic to send to each gateway and via which path. Since wireless channel changes occur fast [1], topology data quickly becomes stale and the shortest path may no longer be optimal just a few seconds later. A better approach is to let the network decide both the path and the gateway on-the-fly as the packet traverses the network, a service we call *dynamic anycasting*. The network may not necessarily deliver the packet to the closest gateway, but rather to one that is available at the moment. Plasma anypath routing empowers the network with such capability.

A. Overview

Plasma anypath routing is inspired by plasma physics. More precisely, our inspiration comes from a plasma lamp, composed of a glass sphere filled with a mixture of gases at a low pressure and an inner orb serving as an electrode [9]. In a plasma lamp, filaments irradiate from the inner electrode towards the outer glass insulator. However, slight variations in the gas temperature may create a region of higher conductivity. Since electric current always takes the path of least resistance, current flows through these high-conductivity regions. As the gas temperature changes, other regions become more conductive and the current will now flow through these regions instead. In our analogy, the inner electrode is the source node and the outer glass sphere represents the set of Internet gateways. Additionally, the electrons in the current denote the packets, the plasma filaments represent the paths taken by the packets, and the high-conductivity regions are the areas with a low interference, low noise, and little multipath fading.

In practical terms, a packet in plasma anypath routing is forwarded to any node within a group of destinations. We refer to this group as the *destination set* and we usually use S to represent it throughout the paper. For wireless mesh networks, the destination set refers to the group of gateways. In plasma anypath routing, a node must keep both a forwarding set J and a transmission rate r for each destination set S. Since each node uses a set of next hops to forward packets, the source node is now connected to the several gateways through a single-source multi-sink directed acyclic graph (DAG), with each node transmitting at a selected rate. Figure 3 depicts this scenario, where each node uses a forwarding set and a bit rate to reach the set of gateways. We define this union of paths, with nodes potentially transmitting at different rates and leading to a subset of gateways as a *plasma anypath*. The multiple bold arrows leaving each node represent the forwarding set. At each hop, only one of the nodes in the set forwards the packet on. As a result, each packet from s traverses one of the available paths to reach either d_1, d_2, or d_3. One of these paths is shown with dashed lines, where a different dash length is used for each bit rate. Succeeding packets, however, may take different paths, with other transmission rates along the way, ultimately reaching different gateways.

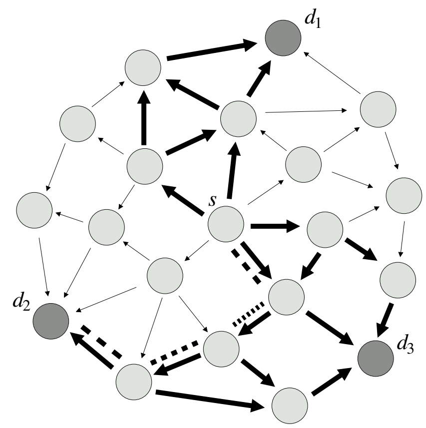

Figure 3. A plasma anypath composed of the union of 14 paths from the source s to the destinations d_1, d_2, and d_3. A packet traverses one of those paths and reaches one of the three gateways, such as the path shown in dashed lines. The different dash lengths represent the different bit rates used at each hop. The routing mechanism is inspired by a plasma lamp, where an inner electrode irradiates filaments towards the outer glass sphere.

In plasma anypath routing, a source node does not know beforehand which gateway will receive its packet. By choosing the forwarding set, the source just selects the potential gateways for a given packet, but it can not guarantee its delivery to a specific one. The final gateway is determined on-the-fly as the packet traverses the network. At each hop, a mesh router selects a single neighbor to keep forwarding the packet. The aggregate of these local decisions results in a global decision for both the path and the gateway of that particular packet. If we consider several packets, we can see that each gateway receives a fraction of the total traffic. Taking advantage of the path and gateway diversity results in an inherent ability of automatically distributing the load among the gateways.

B. Plasma Anypath Cost

The cost of a plasma anypath is calculated in almost the same way as in multirate anypath routing. Each node i has a minimum cost $D_i = \min_{r \in R} D_i^{(r)}$, measured as the average end-to-end transmission time to reach any of the gateways. The cost at a given rate $D_i^{(r)} = d_{iJ}^{(r)} + D_J^{(r)}$ has the same components as before, where $d_{iJ}^{(r)}$ represents the hyperlink cost from node i to the forwarding set J when i transmits at rate r, but $D_J^{(r)}$ now represents the average cost from J to *any* of the selected gateways.

Consider the network in Figure 4, where the weights are the link delivery ratios. For simplicity, let us assume that every node transmits at 1 Mbps and uses 1500-byte packets. The cost via J in Figure 4(a) is calculated as

$$\begin{aligned} D_i &= d_{iJ} + D_J \\ &= \frac{12,000 \text{ bits/1 Mbps}}{1-(1-0.3)(1-0.2)} + \frac{(0.3)13.3 + (0.7)(0.2)15}{1-(1-0.3)(1-0.2)} \\ &= 27.3 + 13.8 = 41.1 \text{ ms}. \end{aligned} \quad (6)$$

One would expect that adding more gateways to the plasma anypath is always beneficial because it provides a higher gateway and path diversity. However, the transmission time can severely increase with more gateways, as shown in Figure 4(b). The cost via J' is $D'_i = d_{iJ'} + D_{J'} = 12.1 + 72.8 = 84.9$ ms, which is more than 2x higher than the cost in Figure 4(a). Since both nodes j and k have a low delivery ratio to d_2 (i.e., 10%), they need to retransmit the packet several times on average to reach d_2. Therefore, when either j or k receives the packet and both nodes in J do not, it is cheaper to retransmit the packet to one of the nodes in J than to take the longer path via j or k. An interesting research question is then to define which gateways a given nodes should use.

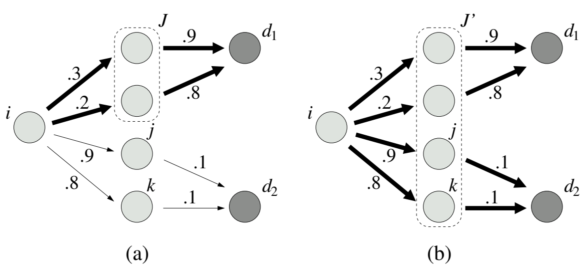

Figure 4. A plasma anypath cost calculation example. The weights are the link delivery ratios. The cost of the plasma anypath in (a) is lower than the cost in (b), even though we have one more gateway available in (b).

An interesting property of the shortest plasma anypath is that it always has a lower or equal cost than the shortest multirate anypath to each individual gateway. This is true because plasma anypath routing is a direct generalization of multirate anypath routing. Among all possible plasma anypaths, we also have the shortest multirate anypaths from the source to each individual gateway. As a result, plasma anypath routing will only choose to use multiple gateways if it is cheaper. Plasma anypath routing is therefore the routing paradigm that provides the *lowest cost* among single-path routing, and both single-rate and multirate anypath routing.

We now address the problem of finding the forwarding set and the transmission rate that provide the least cost for each node to communicate with the best subset of gateways. We call this the *shortest plasma anypath problem*.

C. Routing Algorithm

We introduce a distributed algorithm that finds the shortest plasma anypath from every node to the destination set S. We call it the Plasma Bellman-Ford (PBF) algorithm, since it generalizes the Bellman-Ford algorithm for plasma anypath routing. The algorithm takes only local information as input and works in iterations. At iteration t, each node i updates three variables:

- D_i^t: the cost from i to the destination set S at iteration t;
- F_i^t: the forwarding set i uses to reach S at iteration t;
- T_i^t: the transmit rate i uses to reach S at iteration t.

Initially, we set $D_i^0 = \infty$, $F_i^0 = \emptyset$, $T_i^0 = \text{NIL}$ for every node, except for the nodes in the destination set S. For every node $s \in S$, we have $D_s^0 = 0$. At iteration t, a node updates its cost D_i^t considering the cost of each neighbor at $t-1$, that is, considering D_j^{t-1} for every neighbor j. Before updating the main variables, each node calculates its individual cost when transmitting at each rate $r \in R$ according to

$$D_i^{(r)} = \min_{J \in N_i} d_{iJ}^{(r)} + D_J^{(r)}, \quad (7)$$

where $D_J^{(r)}$ is the remaining cost considering the cost of each neighbor in the previous iteration, as shown below

$$D_J^{(r)} = \sum_{j \in J} w_{ij}^{(r)} D_j^{t-1}, \quad (8)$$

and N_i is the set of optimal forwarding sets for node i. We showed in [7] that, given a set of neighbors with distances $D_1 \leq D_2 \leq \ldots \leq D_n$, the optimal forwarding set is always one of $\{1\}, \{1,2\}, \{1,2,3\}, \ldots, \{1,2,\ldots,n\}$. Forwarding sets with gaps between the neighbors, such as {2,3} or {1,4}, can never yield the lowest cost. As a result, we do not need the test the power set derived from the set of neighbors and therefore we have

$$N_i = \{\{1\}, \{1,2\}, \{1,2,3\}, \ldots, \{1,2,\ldots,n\}\}, \quad (9)$$

which is linear in the number of neighbors. We also store the selected forwarding set used for each rate as

$$F_i^{(r)} = \arg\min_{J \in N_i} d_{iJ}^{(r)} + D_J^{(r)}. \quad (10)$$

Once we have the lowest cost for every rate, the main variables are updated with

$$\begin{aligned} D_i^t &= \min_{r \in R} D_i^{(r)}, \\ T_i^t &= \arg\min_{r \in R} D_i^{(r)}, \\ F_i^t &= F_i^{(T_i^t)}. \end{aligned} \quad (11)$$

The algorithm terminates when $D_i^t = D_i^{t-1}$ for very node i. At termination, each node knows the forwarding set and transmission rate that provide the lowest cost to the optimal destination subset, which is a subset of the destination set S.

Figure 5 depicts the step-by-step execution of the PBF algorithm. For simplicity, we restrict the transmission rate to 11 Mbps and assume that 1500-byte packets are used. The weights represent the link delivery ratios when transmitting at 11 Mbps. Figure 5(a) shows the graph just after the initialization phase. Figures 5(b)–5(d) show each iteration of the algorithm. At each step, the values inside each node represents their current cost estimate D_i and the arrows in boldface represent the shortest plasma anypath to the gateways.

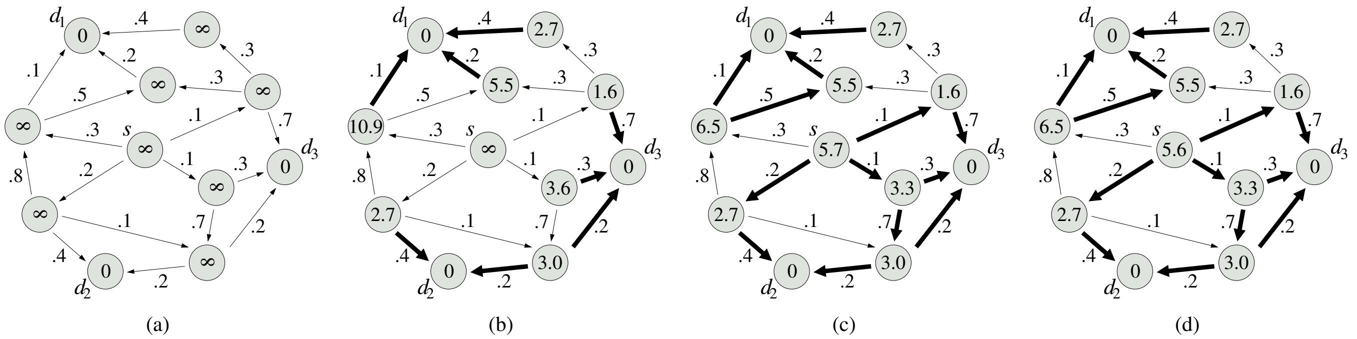

Figure 5. The execution of the Plasma Bellman-Ford (PBF) algorithm from every node to the destination set $\{d_1, d_2, d_3\}$. The weight on each link represent its delivery ratio at 11 Mbps. (a) The situation just after the initialization. (b)-(d) The situation after each successive iteration of the algorithm. Part (d) shows the situation after the last iteration.

The costs of each node are updated considering the costs of the neighbor nodes in the previous step. Figure 5(d) shows the result of the algorithm after the last iteration.

We can make two interesting observations from Figure 5(d). First, the source node s only uses d_2 and d_3 in the optimal gateway subset, but not d_1. Adding d_1 to this subset results in a higher overall cost, thus d_1 is naturally excluded. Second, compared to the single-gateway approach, plasma anypath routing provides a lower cost. The final average cost from s to $\{d_2, d_3\}$ is 6.2 ms. If we had used a single gateway instead, s would have a cost of 8.9 ms to d_1, 9.1 ms to d_2, and 8.3 ms to d_3, which are 25%-32% higher. This additional cost results in a longer medium time, and therefore in a lower bandwidth utilization and throughput.

The running time of the PBF algorithm depends on how the minimization in (7) is implemented. We run the algorithm on a graph $G = (V, E)$ composed of a set V of vertices and a set E of edges or links. The initialization phase takes $O(V)$ time. Assuming the forwarding sets of N_i are tested in the order $\{1\}, \{1, 2\}, \ldots, \{1, 2, \ldots, n\}$, the running time to calculate the cost $d_{iJ}^{(r)} + D_J^{(r)}$ for each forwarding set J takes just $O(1)$, as we showed in [7]. The minimization in (7) checks each pair of link and rate once, with an aggregate time of $O(ER)$. We need at most $|V - 1|$ iterations to converge, for a total cost of $O(V + VER)$, which basically reduces to $O(VER)$. This is the same complexity of the Bellman-Ford algorithm for multiple transmission rates. Therefore, even with exponentially many possibilities of gateways and paths to choose from, finding the shortest plasma anypath is no harder than the shortest-path problem.

We now prove the optimality of the algorithm. Let $\delta_i^{(r)}$ be the cost of the shortest plasma anypath from i to its optimal destination subset, when i transmits at a fixed rate $r \in R$. We use δ_i without the indicated rate to represent the minimum cost among all rates, that is, $\delta_i = \min_{r \in R} \delta_i^{(r)}$. The following theorem proofs that we have $D_i = \delta_i$ for every node at the end of the algorithm.

Theorem 1: Optimality of the algorithm.

Let $G = (V, E)$ be a weighted, directed graph and S be the destination set. After running the Plasma Bellman-Ford (PBF) algorithm on G, we have $D_i = \delta_i$ for every node $i \in V$.

Proof: We prove this theorem in two parts. In the first part, we prove that the algorithm finds the shortest multirate anypath for the single-gateway case. In the second part, we generalize it for the multi-gateway case.

We prove the single-gateway case by induction on t, the iteration number. Let d be the destination (i.e., the gateway) and h_i be the number of hops of the longest path from i to d. We show that, after the t-th iteration, we have $D_i = \delta_i$ for every node with $h_i \leq t$. Intuitively, the algorithm works from the destination backwards to the source in an expanding-ring fashion, settling at each iteration the nodes one hop further away from the destination. Since in a graph with $|V|$ nodes we can not have paths with more than $|V| - 1$ links, after $|V| - 1$ iterations we are guaranteed to have $D_i = \delta_i$ for every node $i \in V$. The induction proof now follows.

Basis. For $t = 0$, the only node with $h_i \leq 0$ is the destination d itself. We have from the initialization that $D_d = \delta_d = 0$.

Inductive step. Assuming that after the t-th iteration we have $D_i = \delta_i$ for every node with $h_i \leq t$, we want to show that after the $(t+1)$-th iteration we have $D_i = \delta_i$ for every node with $h_i \leq t + 1$. At the $(t+1)$-th iteration, a node i with $h_i = t + 1$ calculates $D_i^{(r)}$ in (7) after checking every candidate forwarding set in N_i, including the optimal forwarding set. Since $h_i = t + 1$, every neighbor j in the optimal forwarding set must necessarily have $h_j \leq t$ and we know from the induction hypothesis that $D_j = \delta_j$. As a result, after the minimization, we must have $D_i^{(r)} = \delta_i^{(r)}$ for every rate $r \in R$. Therefore, after selecting the best rate from (11) we now have $D_i = \delta_i$.

For the multi-gateway case, we consider a destination set $S = \{d_1, d_2, \ldots, d_n\}$ instead of a single destination. From the original graph $G = (V, E)$, we construct an extended graph $G' = (V \cup \{d\}, E \cup \{(d_1, d), \ldots, (d_n, d)\})$, with a supernode d to which every destination $d_i \in S$ connects. The weight of the links $(d_1, d), (d_2, d), \ldots, (d_n, d)$ is set to zero.

We claim that the shortest multirate anypath from a node i to d in the new graph G' is also the shortest plasma anypath from i to its destination subset. The shortest multirate anypath is composed of the set of paths that minimize the cost from i to d. All of these paths from i must pass through at least one but not necessarily all of the destinations $d_1, d_2, \ldots, d_n$

to reach d. In fact, the shortest multirate anypath uses the subset of destinations that provides the lowest cost from i to d, choosing the best forwarding sets and transmission rates at each node. This is precisely the definition of the shortest plasma anypath. ∎

D. Load Balancing

In our plasma lamp analogy, when a hand is placed near the lamp, it changes the high-frequency electric field, causing filaments to extend from the inner electrode to the point of contact [9]. An outsider can then drive these filaments around the sphere at will. Ideally, network operators would benefit from an analogous scheme, where the load could be easily moved from one gateway to the next with little effort. This feature is particularly useful when there is a significant mismatch between the received wireless load and the available wired capacity. In this situation, a gateway is able to receive most of the packets coming on its wireless interface, but is unable to forward them all because of its Internet bandwidth is limited, causing many packet drops.

The key idea of the proposed load balancing scheme for plasma anypath routing is to initially assign a non-zero cost for each gateway. That is, in the initialization phase, we assign $D_s = w_s$ for each node $s \in S$. Assigning a non-zero initial cost to loaded gateways leads to a back-pressure effect, as shown in Figure 6. Figure 6(a) shows the result of the PBF algorithm from our previous example, for convenience. In this case, we assume that the traffic received at each gateway is comparable and therefore the load is balanced. However, assume that after some time more mobile clients associate with the mesh nodes near gateways d_2 and d_3, leading to a significant increase in the load at these gateways whereas d_1 is under-utilized. In this case, it makes sense to shift some load from d_2 and d_3 to d_1. Figure 6(b) depicts this scenario, where increasing the weights w_1, w_2, and w_3 to 1, 9 and 7, respectively, reroutes the traffic from several nodes to d_1. The values inside each node i represents the final cost D_i to the destination set. However, the values in Figure 6(b) do not represent the actual end-to-end transmission time anymore, since they are increased by the weight of the gateways.

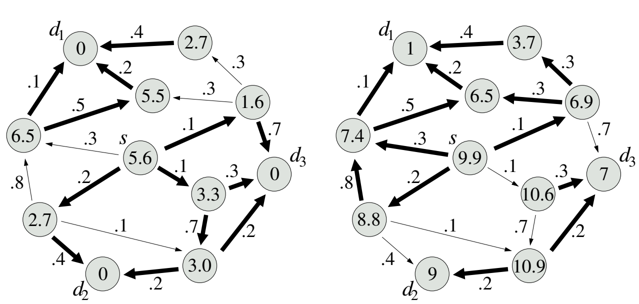

Figure 6. A load balancing scenario. (a) The plasma anypath assuming the same load at the gateways. (b) As gateways d_2 and d_3 receive more traffic, their weights w_2 and w_3 increase, shifting some load to d_1.

IV. Evaluation

We evaluated the proposed plasma routing algorithm with an 18-node 802.11b indoor testbed. We use the ceiling testbed to measure the delivery ratio of each link at different transmission rates. For that purpose, each node broadcasts one thousand 1500-byte packets and later we collect the number received packets at neighbor nodes. We repeat this process for 1, 2, 5.5, and 11 Mbps to have a link estimate for each transmission rate. We use the Click toolkit and a modified version of the MORE software package [4] for the data collection. A complete description of our testbed is given in [7]. Once we have the traces, we use the proposed algorithm to compare the effectiveness of plasma to multirate anypath routing.

Figure 7(a) depicts the cumulative distribution function of the link delivery ratio for each rate. As the rate increases, less neighbors are available and thus path diversity decreases. The intersection of the dotted horizontal line with each curve represents the median. We have only 31 links at 1 Mbps with a delivery ratio lower than 50%. For other rates, this number increases to 73, 85, and 135 for 2, 5, and 11 Mbps, respectively. This imposes a tradeoff for plasma routing, where a higher rate not only reduces the delivery ratio, but also the number of neighbors a node can include in its forwarding set. Our algorithm explores this tradeoff and selects the optimal forwarding set and rate to reach the best gateway subset.

Plasma generalizes multirate anypath routing to anycasting, allowing a source to exploit multiple gateways. As an advantage, the shortest plasma anypath *always* has an equal or lower cost than the shortest multirate anypath to each destination individually. Otherwise, we would have a contradiction, since we can find another plasma anypath (i.e., the single-gateway plasma anypath) with a lower cost to the destination set. To quantify the improvement, we define the gain of plasma anypath routing for a pair (i, S) of source node i and destination set S as the ratio between the shortest plasma anypath from i to S and the shortest multirate anypath from i to the lowest-cost gateway in S.

Figure 7(b) depicts the gain in the end-to-end transmission time of plasma over multirate anypath routing. In the figure, we plot the 2,500 pairs with the maximum gain to show the benefits of plasma anypath routing. The pairs of source node and destination set are placed in order from largest to smallest (i.e., in rank order). The points of each curve are sorted separately and, therefore, the gains of a given x-value are not necessarily from the same pair. For two gateways, we have a maximum gain of 31% while for four gateways this number increases to 54%. For six and eight gateways, the maximum gain is 64% in both cases.

Figure 7(c) shows the cumulative distribution function (CDF) of the expected end-to-end transmission time for every pair (i, S) of source node i and destination set S. As the size of the destination set increases, each node has new gateways to which it may send traffic. Therefore, path diversity increases and packets take potentially closer routes. For one gateway, we have a maximum time of 9.47 ms, with an average of

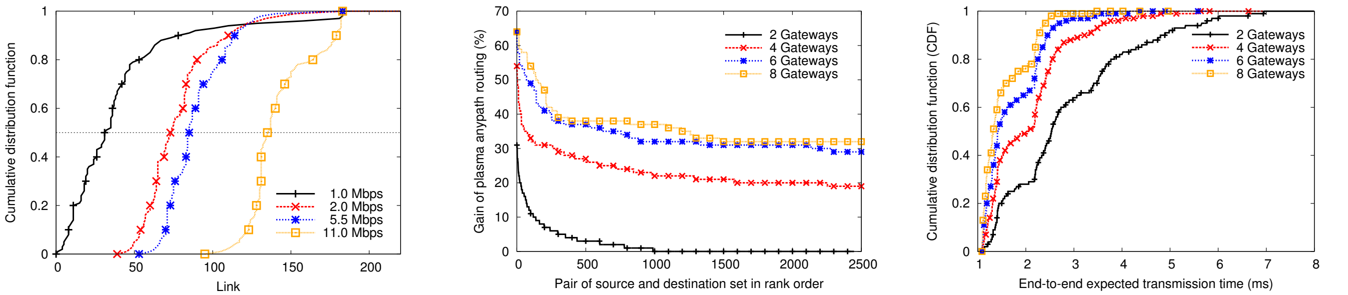

(a) (b) (c)

Figure 7. (a) The cumulative distribution function of the testbed links for each transmission rate. (b) The gain of plasma over multirate anypath routing in the end-to-end expected transmission time. (c) The cumulative distribution function (CDF) of the transmission time for each pair of source node and destination set, for a different number of gateways.

4.6 ms. With two gateways, the maximum decreases to 7.99 ms and the average time is 3.6 ms. For three and four gateways, the maximum end-to-end transmission time is reduced to 7.67 and 6.93 ms, with an average of 3.17 and 2.90 ms, respectively. As we increase the destination set, the expected transmission time approaches 1.1 ms, which is the time a 1500-byte packet takes to be successfully transmitted at 11 Mbps without any retransmissions. More precisely, as more gateways are deployed, nodes are able to reach at least one of them at 11 Mbps in just one hop, without retransmissions. This is in fact the best possible scenario for an 802.11b network, since nodes transmit just once and at the highest speed.

Figure 8(a) depicts the rate selection as we increase the number of gateways in the network. Specifically, we plot the optimal transmission rates selected by each node to reach every possible destination set. As we increase the number of gateways, we see an interesting behavior of nodes migrating from lower to higher bit rates. For instance, initially we have 46.3% of the nodes transmitting at 5.5 Mbps, and 53.7% using 11 Mbps. When we have four gateways, only 29.7% transmit at 5.5 Mbps while 70.3% prefer 11 Mbps. Finally, with eight gateways, 15.4% of the nodes chooses 5.5 Mbps as its optimal transmission rate and 84.6% select 11 Mbps. This is consistent with Figure 7(c) and it is mainly due to the higher gateway density. With more and more gateways distributed across the network, more direct neighbors become gateways. As a result, all of them are included in the forwarding set and it is more likely that a high-rate transmission is successful.

Figure 8(b) depicts the load balancing mechanism in effect. We select two nodes (GW1 and GW2) that are at opposite ends of the testbed to serve as our two gateways. We assign an initial value of zero to GW1 and a value of five to GW2, so that GW1 receives a higher load. We then calculate the shortest plasma anypaths from every node to both of these gateways. We assume that each node injects a unitary load into the network. Figure 8(b) shows the load at each gateway as we increase the weight w_1 of GW1 from zero to one. Increasing w_1 results in shifting the load from GW1 to GW2, as seen in the figure. Initially, GW1 receives 68% of the load whereas GW2 receives only 32%. As we increase w_1, we see the load proportionally moving from GW1 to GW2. When w_1 is roughly 0.3, the load is equally balanced between the two gateways. As we keep increasing w_1, the situation is reversed and GW2 ends up receiving more load. An interesting effect observed in Figure 8(b) is that the load variations occur in a series of discrete steps. This occurs mainly because the weight w_1 must increase up to a point where the routes leading to GW1 become too costly. At this stage, it is shorter for some nodes to switch and use the available routes to GW2 instead.

An interesting question is to quantify how robust plasma routing is in the absence of up-to-date topology information. Figure 8(c) shows this result. We keep the same two gateways as before and calculate the shortest plasma anypath from every node to them. We then fix the forwarding sets and transmission rates of each node in order to compare it over time. We collect topological data every 20 minutes during a 24-hour period. Figure 8(c) depicts the relative difference between the optimal end-to-end expected transmission time and the one we get with the forwarding set and bit rate fixed from the first topology. This gives us an insight of how much we lose by not using up-to-date topology information during routing. We only show the four nodes with the highest differences. Node N1 was the one with the highest difference, showing a maximum of 71.3% approximately 7 hours later. This means that the cost of N1 to the destination set was 71.3% higher than its optimal cost at that particular time. The difference is due to the usage of the forwarding set and transmission rate that were optimal 7 hours before. Its average, however, is only 4.1%. Node N4 was the second one with the highest difference, showing a peak of 25.3% approximately 9 hours later, but with an average of just 6.8%. Nodes N2 and N3 had a maximum difference of 13.8% and 15.5%, respectively. Their averages were, however, much lower, just 2.4% for N2 and 4.2% for N3. For other nodes, the average difference was no higher than 1.5%. Plasma anypath routing therefore does not require a high routing protocol overhead, as long as we use routes which are just a little suboptimal.

V. Related Work

Biswas and Morris [3] designed and implemented ExOR, the first opportunistic routing protocol for wireless multihop

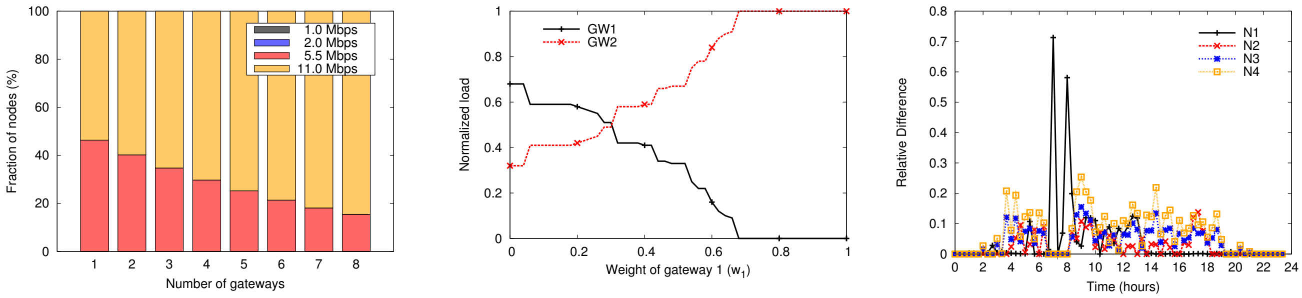

(a) (b) (c)

Figure 8. (a) The transmission rate distribution as a function of the number of gateways. Higher transmission rates are used as the number of gateways increases. (b) Load balancing in effect. Increasing w_1 shifts the load from GW1 to GW2. (c) The robustness of plasma routing over time.

networks. The authors show that opportunistic routing increases throughput by 2x to 4x when compared to single-path routing. Chachulski *et al.* [4] introduce MORE, a routing protocol which uses both opportunistic routing and network coding to further increase ExOR throughput. Neither MORE nor ExOR, however, take full advantage of the multiple transmission rates in 802.11 and of the gateway diversity in wireless mesh networks, which improves the performance even further. An implementation of plasma routing and network coding is part of our future plans.

Dubois-Ferrière *et al.* [5] introduced a shortest anypath algorithm that finds optimal forwarding sets. The authors generalize the well-known Bellman-Ford algorithm for anypath routing and prove its optimality. In [7], we generalize anypath routing for multiple rates. With multiple transmission rates, there is a tradeoff between a higher throughput and lower neighbor diversity. We propose an optimal algorithm capable of selecting both the forwarding set and the bit rate that minimize the cost of every node to the destination. Please refer to [6], [7] for other works addressing opportunistic routing. All proposed algorithms, however, are based on a unicast delivery model and do not consider the effect of anycasting with anypath routing. To our knowledge, we are the first to consider this problem.

Related to wireless anycast routing, Lakshmana *et al.* [14] and references therein share our ideas on each mesh node being simultaneously associated with multiple gateways. Lenders *et al.* [15] propose an anycast routing strategy for wireless ad hoc networks, where nodes can forward packets to an area with a higher gateway density. However, both works do not consider anypath along with anycast routing. The source selects the gateway beforehand for each packet, not allowing the network to take advantage of the available gateways and paths at the moment the packet traverses the network. In contrast, we pre-select a variety of paths to potentially many gateways and the network decides which of those paths the packet should take.

VI. Conclusions

In this paper, we introduced a new routing paradigm for wireless mesh networks. Our key idea is that a mesh node should not be associated with a single gateway for Internet communication. This scenario may lead to both unfairness and under-utilization in the case of hotspots. In plasma anypath routing, the network is responsible for delivering packets to one of the gateways via one of the available paths. We propose a polynomial-time distributed routing algorithm to calculate the forwarding set, transmission rate, and gateway subset that minimizes the cost of every node to the destination set. The algorithm has the same running time as the Bellman-Ford algorithm, and it is therefore suitable for distance-vector routing protocols. We also introduce a load balancing technique that redistributes the load among gateways.

We validate our routing algorithm in an indoor 18-node 802.11b testbed. Our main findings are: (i) plasma anypath routing always outperfoms multirate anypath routing as the number of gateways increases, with a maximum gain of 31% when just two gateways are used and 64% for four gateways; (ii) more nodes select the highest rate (i.e., 11 Mbps) as gateway density increases, resulting in a lower per-packet medium time; (iii) the expected end-to-end transmission time of nodes converges to 1.1 ms, which is the time it takes to transmit a 1500-byte packet at 11 Mbps without retransmissions, and therefore packets are usually delivered at the highest speed at the first try; (iv) it is possible to adjust the load of the gateways by carefully selecting the weights of each node, and (v) plasma routing is robust to wireless link fluctuations; in a one-day period without updating the topology information, the cost of a plasma anypath is on average less than 7% higher than the same cost assuming full knowledge of the topology.

References

[1] M. E. M. Campista, D. G. Passos, P. M. Esposito, I. M. Moraes, C. V. N. de Albuquerque, D. C. M. Saade, M. G. Rubinstein, L. H. M. K. Costa, and O. C. M. B. Duarte, “Routing metrics and protocols for wireless mesh networks,” *IEEE Network*, vol. 22, no. 1, pp. 6–12, Jan.-Feb. 2008.

[2] I. M. Moraes, M. E. M. Campista, L. H. M. K. Costa, O. C. M. B. Duarte, J. L. Duarte, D. G. Passos, C. V. N. de Albuquerque, and M. G. Rubinstein, “On the Impact of User Mobility on Peer-to-Peer Video Streaming,” *IEEE Wireless Communications*, vol. 15, no. 6, pp. 54–62, Dec. 2008.

[3] S. Biswas and R. Morris, “ExOR: Opportunistic Multi-Hop Routing for Wireless Networks,” in *ACM SIGCOMM’05*.

[4] S. Chachulski, “Trading Structure for Randomness in Wireless Opportunistic Routing,” Master’s thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, May 2007.

[5] H. Dubois-Ferrière, “Anypath Routing,” Ph.D. dissertation, Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland, Nov. 2006.
[6] M. Lu and J. Wu, “Opportunistic Routing Algebra and its Applications,” in *IEEE Infocom’09*.
[7] R. Laufer, H. Dubois-Ferrière, and L. Kleinrock, “Multirate Anypath Routing in Wireless Mesh Networks,” in *IEEE Infocom’09*.
[8] B. He, B. Xie, and D. P. Agrawal, “Optimizing deployment of Internet gateway in Wireless Mesh Networks,” *Computer Communications*, vol. 31, no. 7, pp. 1259–1275, May 2008.
[9] Wikipedia, The Free Encyclopedia, “Plasma lamp,” [Online] http://en.wikipedia.org/wiki/Plasma_lamp.
[10] S. Jain and S. R. Das, “Exploiting Path Diversity in the Link Layer in Wireless Ad Hoc Networks,” *Ad Hoc Networks*, vol. 6, no. 5, pp. 805–825, Jul. 2008.
[11] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, “Measurement-Based Models of Delivery and Interference in Static Wireless Networks,” in *ACM SIGCOMM’06*.
[12] D. D. Couto, D. Aguayo, J. Bicket, and R. Morris, “A High-Throughput Path Metric for Multi-Hop Wireless Routing,” in *ACM MobiCom’03*.
[13] R. Draves, J. Padhye, and B. Zill, “Routing in Multi-Radio, Multi-Hop Wireless Mesh Networks,” in *ACM MobiCom’04*, pp. 114–128.
[14] S. Lakshmanan, R. Sivakumar, and K. Sundaresan, “Multi-Gateway Association in Wireless Mesh Networks,” *Ad Hoc Networks*, vol. 7, no. 3, pp. 622–637, May 2009.
[15] V. Lenders, M. May, and B. Plattner, “Density-Based Anycast: A Robust Routing Strategy for Wireless Ad Hoc Networks,” *IEEE/ACM Transactions on Networking*, vol. 16, no. 4, pp. 852–863, Aug. 2008.