Gateway-aware Routing for Wireless Mesh Networks

Prashanth A. K. Acharya, David L. Johnson, Elizabeth M. Belding
Department of Computer Science
University of California - Santa Barbara
(acharya, davidj, ebelding)@cs.ucsb.edu

***Abstract*—Wireless mesh networks (WMNs) provide an attractive method to provide Internet connectivity in developing regions. Traditional mesh routing protocols are designed to find high quality/throughput multihop routes in the network. However, these solutions do not consider constraints imposed by the capacity at the gateway, often the bottleneck in such rural area networks. In this paper, we demonstrate the importance of intelligent choice of gateways in WMNs. We present the design of a new gateway-aware routing metric that picks high throughput routes in the presence of heterogeneous gateways. Our evaluation in simulations as well as on a testbed show significant increase in network throughput.**

I. Introduction

IEEE 802.11-based wireless mesh networks (WMNs) have been used as an economic and convenient technology to provide connectivity to rural areas, especially in developing regions[1], [2], [3], [4]. Even in developed countries, community-wide WMNs have provided Internet connectivity to a large geographic area. We envision the following architecture to use WMNs for Internet access in these areas. Wireless routers are placed in homes to provide in-residence Internet access. A small number of these routers are connected to the Internet and function as gateway nodes for the network. These gateway nodes may use technologies such as DSL, WiMAX, 1xEVDO, cellular, long distance WiFi or satellite as means for capacity injection/backhaul connection. Backhaul connections may be hosted and shared by local businesses or by members of the community.

Traditional mesh routing solutions have focused on simply finding the best route to the gateway to reach the Internet. Two important underlying assumptions of these routing solutions are: 1) all gateway nodes are equally capable in terms of resources such as bandwidth capacity and delay to connect to the Internet; and/or 2) the capacity *bottleneck* is in the wireless multihop portion of the WMN. We now examine each of these assumptions in the context of rural and community WMNs.

In some WMN deployments, all the gateway nodes are similarly provisioned. However, this is not the case in many other scenarios. In rural/developing regions, cost and availability considerations influence the uplink connectivity options that can be used. Another scenario is a neighborhood community WMN wherein volunteers share a portion of their existing Internet connections. These may be DSL, cable-modem, WiMAX or GPRS/EDGE and leads to a heterogeneous mix of gateway uplinks for the WMN. Many Internet connections do not provide guaranteed bandwidth, but instead the capacity varies with network load. Clearly, the assumption of homogeneous gateways does not always hold.

The second assumption in traditional routing is that the capacity of the multihop mesh route is less than that of a gateway uplink. This is more common in a pure IEEE-802.11b network with a raw capacity of 11Mbit/s (and typical multihop capacity of the order of 1Mbit/s [5]). However with the advent of newer 802.11 radio technologies, the capacity of mesh networks has increased by orders of magnitude. (54 Mbit/s for 802.11g or 802.11a and up to 300 Mbit/s for 802.11n). These capacities may be significantly higher than those of the common uplink technologies, e.g. the download capacity for DSL ranges from 1.5Mbit/s to 24Mbit/s, the typical capacity of a WiMAX client is 4Mbit/s [6], a VSAT connection is 1024kbit/s [3]. Therefore, in many scenarios, the capacity of the mesh route may be more than that of the gateway.

The throughput performance of clients connected to a mesh router, intuitively, is influenced by the routing path to the Internet, including the chosen gateway node. Traditional routing solutions, however, have dealt primarily with routing inside the mesh with wireless capacity maximization and interference minimization as prominent objectives. Clearly, in WMNs with varied gateway capacities, the choice of gateway also has a large influence on the performance of the network. Therefore, we believe the correct question to consider in routing is the following: *"at each mesh router, what is the best route to access the Internet?"* This problem statement captures the constraints introduced by the gateway capabilities in addition to the problems of wireless capacity and interference.

In this paper, we propose a gateway-aware mesh routing solution that intelligently selects gateways for each mesh router based on the multihop route in the mesh as well as the capability of the gateway. We develop a new composite routing metric called Gateway-aware Routing Metric (GARM) that captures these aspects of routing in WMNs. We implement our routing metric for the upcoming IEEE-802.11s (mesh networking) standard in a simulator. Further, we evaluate this metric on a testbed network based on an existing routing solution [5]. Our evaluations demonstrate the importance of intelligent gateway selection in order to improve network performance. We show that the GARM metric can increase the overall network throughput, in some cases by 22%, through better utilization of existing resources.

II. Related Work

There is a significant body of research work that focuses on improving the performance of WMNs. These have addressed various aspects of WMN operations including routing, channel assignment, and interference management. However, most of these solutions either consider the network to have a single gateway, or multiple homogeneous gateways.

Multi-gateway systems have been studied in the context of load balancing among the gateways. The Hyacinth architecture considers multiple gateways and suggests gateway selection based on available bandwidth, to achieve load balancing [7]. This architecture requires accurate knowledge of the flow bandwidth requirements at each mesh router, which may be non-trivial to estimate. Nandiraju et al. propose dynamic gateway switching in a WMN to achieve load balancing: the gateway node monitors congestion levels on its uplink, and at the onset of congestion sends a message to an associated mesh router to choose a different gateway [8]. In contrast, our approach is proactive and does not require frequent congestion measurement. Lakshmanan et al. suggest the simultaneous use of multiple gateways by striping packets of a flow to maximize uplink capacity utilization [9]. Their architecture, however, requires a super-gateway that handles packet re-ordering caused by packet striping. Such an architecture may not be possible in rural WMNs because the uplinks may be provided by different ISPs. Other systems, as proposed in [10], formulate the problem as a linear optimization problem. This approach, however, requires global knowledge of gateway and link capacities, and is centralized. On the other hand, our system is distributed and builds upon the existing routing framework.

III. Background

We now outline our assumptions and considerations for the design of a gateway-aware routing protocol in WMNs.

Assumptions: We assume the mesh network consists of a multi-tiered architecture with a mesh backhaul layer responsible for communication among the mesh routers (including the gateway node), and a client access layer that communicates with end-user devices. This architecture is commonly used by several real-world deployments [2], [5]. Our protocol is currently designed for a mesh backhaul layer that uses one 802.11 radio at each node. We assume that only mesh nodes use our routing protocol and select one gateway at a time, i.e., multi-gateway associations are not permitted.

Design Preliminaries: We consider the problem of finding the best route to the Internet from any node, given a choice of multiple gateway nodes and several candidate multi-hop routes to reach these gateways. The gateways are characterized by different capacities, and the candidate routes by different path qualities. A naïve approach would be to include an additional virtual node connected to each gateway with a link capacity corresponding to each gateway's capacity. The problem is now reduced to finding the best route from each mesh node to this virtual node. The uplink is mapped as an additional wireless link and existing routing algorithms are used.

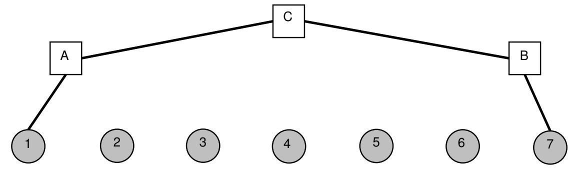

Fig. 1. Simple Line Topology: The gray circles are mesh nodes. Gateway nodes 1 and 7 are connected to routers A and B via Ethernet-like connections. We measure the throughput to node 5.

We conduct simulation experiments using Qualnet [11] to understand the feasibility and/or pitfalls of this simple approach. In particular, we study the effects of gateway capacity on the throughput performance of the mesh network.

We consider the line topology of Figure 1 with seven equidistant mesh nodes. This topology also approximates the decision process in a more complex complete graph where at some point two paths extend to either gateway. In each experiment trial, the distance between the nodes is varied to ensure that only adjacent nodes can communicate directly. All the mesh nodes use the IEEE 802.11s mesh networking extension with HWMP routing protocol to form the mesh topology. The radios operate using the 802.11b/g MAC and PHY standard and bit-rate is fixed uniformly across the network. All other parameters are set to the default values provided by the simulator. For each experiment trial, we use a different bit-rate in order to vary the wireless mesh capacity. Nodes 1 and 7 are the gateway nodes and are connected to the external nodes A and B via wired-links. The capacities of these links are varied to simulate different capacities. Nodes A and B are connected to node C with a high capacity link. Node C forms one end-point of all communication and represents the larger Internet.

In order to study the impact of gateway choice, only one of the uplinks (Link 1A or 7B) is active at any time. We initiate one UDP flow from node C to node 5 and measure the throughput for this flow through each gateway. Table I shows the average of five trials of these experiments.

Gateway Node	Gateway Capacity (Mbps)	Wireless bit-rate (Mbps)	UDP throughput (Mbps)
1	1.5	36	1.443
7	0.5	36	0.493
1	1.5	2	0.367
7	0.5	2	0.458

TABLE I
Throughput with two gateways in a line topology.

We make two key observations from these results. First, using a gateway with a high capacity uplink can lead to higher throughput, even though it may take more hops (or mesh air-time) to communicate with the gateway. Therefore, an intelligent choice of gateway can indeed lead to better network performance. Second, using a lower wireless bit-rate, the throughput while using the closer gateway with lower uplink capacity is higher than that using a high capacity distant gateway. We conclude that the total end-end throughput is dominated by one of either the gateway capacity or the capacity of the mesh path, i.e., one of these acts as a bottleneck for the path and effectively determines the maximum achievable throughput.

The above observations, imply that the protocol should treat the two parts of the path (i.e., multihop wireless path and the gateway uplink) independently. The metric should be able to identify the bottleneck portion of each path, and choose an appropriate gateway. The challenge, therefore, is to design a routing metric that simultaneously considers the gateway capacity as well as the quality of the multihop route from a given node to the gateway. The design of such a metric necessitates a mapping of gateway capacity to the path quality, in order to compare and identify the bottleneck. We next present the design of the GARM metric that enables joint selection of gateways and routes to these gateways.

IV. Design

Based on the observations in Section III, we can consider a gateway-aware route in a WMN to consist of two parts: first, the multihop wireless path; second, the uplink at the gateway node. We present the metric design for these components, and then describe the unified metric.

Metric for the wireless path: Yang et al. identify four fundamental characteristics that a WMN routing metric should capture [12]. The routing metric should increase with the length of path, capture packet loss ratios of the links, consider link capacities, and help reduce interference in the channel shared by the wireless links. Early research on routing for multihop wireless networks commonly used hop count as the routing metric. This metric incorporates only the first characteristic. The Expected Transmission Count (ETX) [13], defined as the expected number of MAC layer transmissions to deliver a packet, only captures the first two characteristics.

The Expected Transmission Time (ETT) metric improves upon ETX by considering the differences in link capacities [14]. The ETT of a link i is defined as the expected duration of a successful MAC layer transmission over the link i. The ETT of a link l is defined by the following relationship:

$$ETT_i = ETX_i \cdot \frac{S}{B_i}$$

where S is the packet size and B_i is the bit-rate for link i. The weight of a path p is the sum of the individual ETTs of the links along the path:

$$mETT = \sum_{l \in p} ETT_l$$

The ETT metric captures the impact of the link capacities on the path performance. For single radio networks, the ETT metric also considers the impact of interference. Therefore, ETT has all the desirable properties for the routing metric of WMN. Routing using the ETT metric has been shown to provide high throughput paths, compared to other previous metrics such as ETX [5], [14]. We therefore choose to use ETT as the metric for the wireless portion of the gateway-aware metric. We also note that the Airtime metric, one of the proposed routing candidate metrics in the IEEE 802.11s standard, is similar to the ETT metric.

Gateway capacity metric: We define the gateway capacity metric $gwETT$ as the time required to transmit a packet of size S on the uplink, and is given by

$$gwETT = ETX_{gw} \cdot \frac{S}{B_{gw}}$$

where B_{gw} is the capacity of the gateway and ETX_{gw} is the expected transmission count for the uplink. In the simple case, we assume that the uplink is a reliable medium or has negligible loss rates, and ETX_{gw} is one. If the uplink is an unreliable medium, we assume that the ETX_{gw} is provided to the routing protocol by some external module (e.g., a link quality measurement tool). Such an ETT-like design of the gateway capacity metric enables a simple mapping and direct comparison to the wireless path ETT-metric.

Combining the two metrics: In Section III, we noted that the throughput of a path is determined by the bottleneck portion of the path. A comparison of the two ETT metrics, $mETT$ and $gwETT$, identifies this bottleneck; the larger of the two indicates that more time is required for transmission on the corresponding portion, and therefore is the bottleneck on the path.

$$M_i = max(mETT, gwETT)$$

This metric, however, is not *isotonic* [12]. Isotonicity is an important property for routing metrics to ensure that efficient algorithms such as Dijkstra can be used to calculate paths with minimum weight. In addition, isotonicity also ensures loop-free forwarding in hop-by-hop routing.

Therefore, we use the following metric for routing to a gateway:

$$GARM = \beta \cdot M_i + (1 - \beta) \cdot (mETT + gwETT)$$

This Gateway-aware Routing Metric has two parts. The first part of the metric accounts for bottleneck capacity. The second factor models the total delay of the path, including the uplink. β is used to control the balance between these two factors. The gateway with the least GARM value is chosen as the default gateway. We note that $mETT$ increases at each hop, and since $gwETT$ is constant for a given gateway, M_i either stays the same or increases at each hop. Therefore $GARM$ is isotonic. This metric requires the propagation of the $gwETT$ metric in addition to the path metric, i.e., each node needs to know of the gateway capacity metric for its routing calculations.

A. Design Discussion

Capacity measurement: One of the inputs to the GARM metric is the gateway capacity. We assume that this value is input to the routing software by an external entity. This capacity input is straight-forward on uplinks with guaranteed/fixed bandwidth, e.g., Ethernet, in a community WMN wherein a volunteer wants to share only a fixed portion of his Internet connection. On uplinks with time-varying capacity, (e.g. wireless uplinks), we assume the presence of capacity-measurement tools such as *pathrate*. In response to a change in uplink capacity, gateway selection may change at the same time as required to propagate the routing information.

Load-sensitivity: GARM does not consider the traffic load on the gateway for gateway selection. Load-based routing leads to routing instability and route flapping [12]. A similar effect would occur for gateway selection as well. However, in networks where the mesh paths are of similar qualities, GARM achieves implicit load-balancing among gateways. Due to the formulation of the $gwETT$ metric and its comparison with $mETT$, GARM associates mesh routers to gateways approximately in proportion to the gateway capacities.

Multi-radio WMNs: The GARM metric is designed for operation on single radio/channel networks. As noted before, operation of the gateway uplink is independent of the wireless path. This behavior can be considered analogous to a wireless link operating on a different channel. Previous research has designed routing protocols such as MR-LQSR with the WCETT metric [14]. We believe such solutions can be extended to be gateway-aware by considering the uplink as a wireless link operating on an orthogonal channel.

V. Implementation and Performance Evaluation

We implement the GARM routing metric in two environments. First, we extend the 802.11s-HWMP implementation in Qualnet 4.5 to include the composite GARM metric. The 802.11s protocol provides a framework for WMN operations including neighbor discovery, link quality estimation, route discovery and route updates. The 802.11s routing beacons that originate from gateway/portal nodes include additional information about the $gwETT$ value corresponding to its uplink capacity. Intermediate routers repeat the $gwETT$ value when they forward the routing beacons, thereby indicating the $gwETT$ value of the path in which the update propagated. Our second implementation of the GARM metric is for the Click-based MIT RoofNet software [5]. Here too, the gateway routing advertisements of the SRCR routing protocol are extended to include the $gwETT$ value, and intermediate routers rebroadcast this information. We use this implementation for our experiments on the UCSB MeshNet testbed (see Sections V-B and V-C).

To evaluate the performance of the GARM metric in WMN routing, we first study the accuracy of the metric in selecting gateways. Subsequently, we characterize the throughput gains obtained by intelligent gateway selection. We compare the performance of our solution with that of routing protocols that use the ETT metric. We evaluate GARM in simulation-based experiments as well on a testbed network. The simulation study allows fine-grained control of network parameters such as topology, link quality, bit-rate, etc. to enable better understanding of the metric. The testbed represents an uncontrolled environment with variable and heterogeneous link qualities, dynamic bit-rate, external interference, etc. and provides a realistic WMN environment.

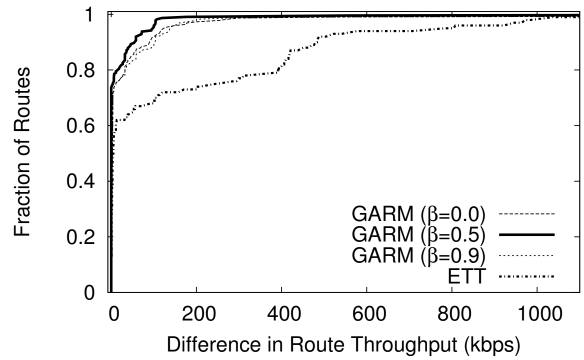

Fig. 2. Accuracy of gateway selection: CDF of difference in throughput using the GARM and ETT routing metrics compared to that using the Oracle approach (best gateway). A perfect routing metric has zero difference.

In the evaluation of GARM we use different combinations of gateway uplink capacities. We choose three gateway capacities that represent the uplink technologies that are used in rural WMNs: 4Mbit/s to represent ADSL, long-distance Wifi, WipLL (Wireless IP Local Loop), WiMAX [1], [6], [15]; 1.5Mbit/s for CDMA450, T1, 1xEVDO [6], [15]; 0.5Mbit/s for a VSAT link or EDGE [3], [15]. The absolute values of the uplink capacities can vary due to the level of service purchased. For example service providers will often provide a guaranteed rate and a maximum burst rate which varies based on the number of users. A possible solution is to advertise a typical mean for a specific day and hour using historic measurements. Calculating actual instantaneous capacity is a non-trivial problem and for now we assume that the gateway capacity is constant for the duration of the experiments. However, it is important to note that there is an order of magnitude difference in their relative capacities. As we demonstrate later, the larger the difference in gateway capacities, the more critical is the role of gateway-aware routing in increasing network performance.

A. Gateway Selection Accuracy

We first use simulation-based experiments to study GARM's accuracy in selecting gateways under different operating environments. Therefore, we conduct experiments with various combinations of uplink and wireless path capacities.

Our experiments use twenty-five 802.11b/g nodes in a 5x5 uniform grid topology. For each experiment, we fix the link bit-rate to be either 2Mbit/s, 11Mbit/s, or 36Mbit/s; all nodes use the same bit-rate. The distance between the nodes is varied accordingly to ensure that only 1-hop neighbors have direct communication. Two diagonally opposite corners of the grid are chosen as gateway nodes. Similar to the topology shown in Figure 1, the gateway nodes are connected to external nodes A and B via wired-links. The capacities of these links are varied. Nodes A and B are connected to node C with a high capacity link. Node C is the end-point of all communication and

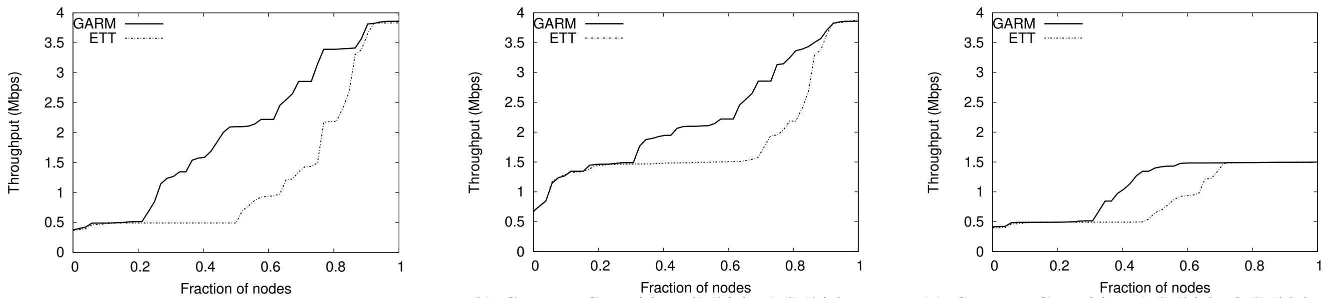

(a) Gateway Capacities: 4Mbit/s, 0.5Mbit/s (b) Gateway Capacities: 4Mbit/s, 1.5Mbit/s (c) Gateway Capacities: 1.5Mbit/s, 0.5Mbit/s

Fig. 3. Distribution of node TCP throughput with different gateway capacities. GARM increases throughput for up to 67% of the nodes. The performance increases as the the difference in gateway capacities grows.

represents the Internet. The two gateway capacities are chosen as a combination of 4Mbit/s, 1.5Mbit/s and 0.5Mbit/s, leading to three distinct uplink capacity pairs. These uplink capacity pairs, together with the link bit-rates, lead to nine different combinations of gateway and wireless path capacities. Through simulations, we verify that GARM selects the appropriate gateway in each scenario.

For each scenario, we perform five trials of the following experiments. We first determine the best gateway for each wireless node in the network. This is achieved by comparing the throughput of a TCP flow from node C to each wireless node, and evaluating the throughput through each gateway. We call this method the Oracle approach because it determines the maximum achievable throughput. Next, we repeat the TCP throughput experiment and let the routing protocol choose the gateway. We record the throughput as well as the choice of gateway. The experiment is repeated for ETT and GARM with β values of 0.0, 0.5 and 0.9. Note that $\beta = 0$ represents the case where the metric does not identify the bottleneck, instead it only considers the sum of $mETT$ and $gwETT$. For every node, we determine whether the routing protocol chooses the best gateway, and calculate the difference in the achieved throughput from that of the Oracle approach.

Figure 2 plots a CDF of the difference in throughput for each route using the GARM and ETT metrics compared to the Oracle approach over all nine capacity combinations. A perfect routing metric would always choose the best gateway (similar to the Oracle approach), and therefore have zero difference in throughput for all routes. From the graph, we see that ETT selects the best gateway for only 60% of the nodes. Further, for 20% of the routes, the loss of throughput due to poor gateway choice is significant (more than 400kbit/s). GARM (with β=0.5), on the other hand, has a very small fraction of routes ($<$3%) wherein the loss of throughput is greater than 150kbit/s. With GARM, 76% of the nodes select the best gateway. For the nodes that do not select the best gateway, we see that the difference in throughput is less than 150kbit/s. This shows that the GARM metric is unable to distinguish between routes that differ by less than 150kbit/s.

The graph shows that the performance of GARM with β values of 0, 0.5, and 0.9 is quite similar. This leads us to believe that GARM is not very sensitive to the value of β. However, we recommend the use of $\beta = 0.5$, where bottleneck capacity and path delay including the uplink are perfectly balanced, since it provides the best results among the three values we tested.

B. Throughput Performance

We next evaluate GARM on the UCSB MeshNet indoor testbed [16]. The testbed consists of 15 nodes deployed in offices on different floors of a building. Each node consists of two Atheros-based 802.11a/b/g radios, and uses the MadWifi driver (v0.9.4) on Linux (kernel v.2.6-15). We use only one of the radios at each node and operate it in 802.11b/g mode. Each node is also connected to a LAN via Ethernet. This provides a control interface to manage the nodes and conduct experiments. This setup also allows us to provision any testbed node to be a gateway node with the Ethernet link as uplink, thereby enabling evaluation of different topologies. We use the MIT RoofNet software for routing in this network. The SRCR routing protocol [5], implemented in Click, uses the ETT metric for routing. Our implementation incorporates GARM in the SRCR protocol. For our experiments, we use a dedicated host machine on the LAN to act as an endpoint for TCP tests and a Web server. In the context of an actual WMN, this dedicated host represents the Internet.

We now characterize the throughput gains achieved at each mesh node with intelligent gateway selection using the GARM metric. We use two gateways in the network, each with different capacity. The capacity of the gateways is again chosen among 4Mbit/s, 1.5Mbit/s and 0.5Mbit/s. We use the `tc` traffic control utility to limit the gateway capacities to these values. The two gateways are randomly chosen from nodes in the network. The selection of two random gateways with a given capacity combination (e.g., 4Mbit/s and 1.5Mbit/s) constitutes a network topology for the experiment. With each topology, we perform the following. For each node in the network, we use the `nuttcp` throughput measurement tool to measure the throughput of a 3-minute TCP stream from the dedicated LAN host to the node. This represents this maximum throughput at the node when downloading a file from the Internet. Only one stream is active at a time. We record the average throughput of each node based on three trials. We repeat the experiment for

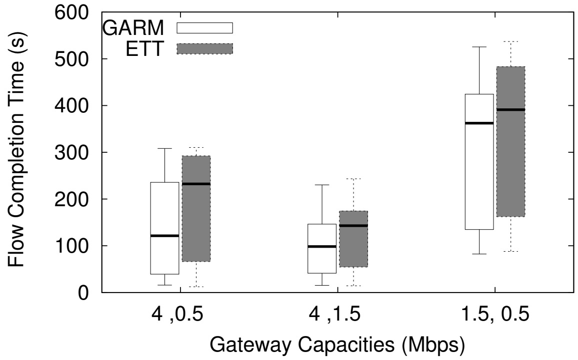

Fig. 4. Distribution of time required to download a 5 MByte file in the UCSB MeshNet.

five different topologies, for both the GARM and ETT routing metrics.

We plot the distribution (CDF) of individual node throughputs for these testbed experiments in Figure 3. Figures 3(a), (b) and (c) show the distribution of throughputs when the capacities of the gateway is 4Mbit/s & 0.5Mbit/s, 4Mbit/s & 1.5Mbit/s, and 1.5Mbit/s & 0.5Mbit/s, respectively. The percentage of nodes that obtain better throughput with GARM are 67%, 53% and 36%, respectively. The average increase in throughput (and percentage increase in throughput) for these scenarios is 679kbit/s (58%), 403kbit/s (24%), and 194kbit/s (22%) respectively. Although GARM provides better throughput for many nodes in all three scenarios, we note that the absolute gain and the number of affected nodes is influenced by the difference in the gateway capacities. The larger the difference, the larger the fraction of nodes affected and greater the throughput improvement.

C. Overall Network Performance

Having evaluated the performance gains of individual nodes in the network, we now characterize the overall network throughput performance. Each node in the network simultaneously downloads a five megabyte file from the dedicated LAN host. At each node, the time required for completion of the download is measured. The experiment is repeated with different gateway capacity combinations and for five different topologies.

Figure 4 shows a box-and-whiskers plot depicting the quartiles of the distribution of flow completion time for the GARM and ETT metrics[1]. The median values are shown in the box and the whiskers represent the minimum and maximum values. From the graph, we observe that GARM reduces the median download time in each case. Further, the distribution of the lower and the median quartiles is skewed towards lower flow completion times, implying that a large number of flows complete sooner when using GARM than when using ETT. This effect is most significant when the difference between the gateway capacities is large. For example, with gateway capacities of 4Mbit/s and 0.5Mbit/s, the average completion time for GARM and ETT are 146.9s and 179.3s, respectively. This represents a 22% increase in average network throughput, through intelligent gateway selection that results in better utilization of network resources.

[1]We choose the box-and-whiskers representation instead of average and standard deviation because the distribution is approximately bimodal; this is because the flow completion time in many cases is limited by the capacity of one of the two gateways.

VI. Conclusion

Wireless mesh networks in developing regions are connected to the Internet with a variety of uplink technologies, each with different capacities. In this paper, we have shown that the selection of gateways for routing plays a very important role in determining the performance of the network. We presented the design of the GARM routing metric that can effectively choose the best gateway for each mesh node. Through performance evaluations we showed that intelligent gateway selection can increase the throughput of nodes in the WMN, and increase the effective capacity of the network. In the future, we wish to extend this framework to incorporate aspects such as the different costs of gateway bandwidth, power availability at each gateway and application-aware gateway selection.

References

[1] S. Surana, R. Patra, S. Nedevschi, M. Ramos, L. Subramanian, Y. Ben-David, and E. Brewer, "Beyond Pilots: Keeping Rural Wireless Networks Alive," in *Proc. of NSDI*, San Francisco, CA, 2008.

[2] P. Bhagwat, B. Raman, and D. Sanghi, "Turning 802.11 Inside-Out," in *Proc. of HotNets*, Cambridge, MA, Nov. 2003.

[3] (2009, May) LinkNet. [Online]. Available: http://link.net.zm/

[4] D. Johnson, "Evaluation of a single radio rural mesh network in South Africa," in *ICTD*, Bangalore, India, Dec. 2007.

[5] J.Bicket, D.Aguayo, S.Biswas, and R.Morris, "Architecture and Evaluation of an Unplanned 802.11b Mesh Network," in *Proc. of MobiCom*, Cologne, Germany, Aug 2005.

[6] S. Mishra, J. Hwang, D. Filippini, R. Moazzami, L. Subramanian, and T. Du, "Economic Analysis of Networking Technologies for Rural Developing Regions," *Lecture Notes in Computer Science*, vol. 3828, pp. 184–194, Nov. 2005.

[7] A. Raniwala and T. Chiueh, "Architecture and Algorithms for an IEEE 802.11-Based Multi-Channel Wireless Mesh Network," in *Proc. of INFOCOM*, Miami, FL, Mar. 2005.

[8] D. Nandiraju, L. Santhanam, N. Nandiraju, and D. Agrawal, "Achieving Load Balancing in Wireless Mesh Networks Through Multiple Gateways," in *Proc. of MASS*, Vancouver, Canada, Oct. 2006.

[9] 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.

[10] S. Tajima, T. Higashino, N. Funabiki, and S. Yoshida, "An Internet Gateway Access-Point Selection Problem for Wireless Infrastructure Mesh Networks," in *Proc. of MDM*, Nara, Japan, May 2006.

[11] (2009, May) Qualnet Network Simulator, Version 4.5. [Online]. Available: http://www.scalable-networks.com

[12] Y. Yang, J. Wang, and R. Kravets, "Designing Routing Metrics for Mesh Networks," in *Proc. of WiMESH*, Santa Clara, CA, Sep. 2005.

[13] D. S. J. De Couto, D. Aguayo, J. Bicket, and R. Morris, "A High-Throughput Path Metric for Multi-hop Wireless Routing," in *Proc. of MobiCom*, San Diego, CA, Oct 2003.

[14] R. Draves, J. Padhye, and B. Zill, "Routing in Multi-radio, Multi-hop Wireless Mesh Networks," in *Proc. of MobiCom*, Philadelphia, PA, Sep 2004.

[15] (2009, May). [Online]. Available: http://en.wikipedia.org/wiki/List_of_device_bandwidths

[16] H. Lundgren, K. Ramachandran, E. Belding-Royer, K. Almeroth, M. Benny, A. Hewatt, A. Touma, and A. Jardosh, "Experiences From Building and Using the UCSB MeshNet Testbed," *IEEE Wireless Communications Magazine*, Apr. 2006.