

Symmetrical Fairness in Infrastructure Access in Multi-Hop Wireless Networks

Saumitra M. Das

Himabindu Pucha

Y. Charlie Hu

School of Electrical and Computer Engineering
Center for Wireless Systems and Applications
Purdue University
West Lafayette, IN 47907
Email: {smdas, hpucha, ychu}@purdue.edu

Abstract

In this paper, we study the problem of providing fairness in multi-hop wireless infrastructure access. In such networks, it is well known that the use of current media access and transport protocols can result in severe unfairness and even starvation for flows originated from different numbers of hops away from a wired infrastructure point or gateway. In this paper, we study a different type of fairness that exists in such networks – flows initiated by nodes that are similar numbers of hops away from the gateway can experience significant unfairness, and such unfairness exists even for perfectly symmetrical node distribution and channel conditions. We denote such fairness as symmetrical fairness. We first provide a framework to characterize and measure symmetrical fairness. We then perform an extensive set of simulation experiments to quantify the causes of symmetrical unfairness. Finally, we develop and study a distributed routing algorithm that significantly improves the symmetrical fairness.

1 Introduction

Providing mobile nodes with access to infrastructure is an important research challenge. We define infrastructure to be entities in a wireless network that nodes want to connect to for various reasons ranging from gaining Internet access, backing up data, reporting sensor readings, or sending status reports to enable command and control. Such infrastructure is likely to be sparsely deployed in wireless networks. For example consider infrastructure points that provide access to the wired Internet by acting as gateways for mobile nodes (possibly providing services like NAT). These gateways with access to the Internet may not be available ubiquitously. Thus a node may experience frequent

disconnections and reconnections or even total loss of service while being mobile. However, it is possible for mobile nodes to use other mobile nodes as forwarders in order to reach an Internet access gateway that is multiple hops away. Thus in such networks, each mobile node that is not in direct transmission range of the gateway uses a multi-hop routing protocol to reach the gateway.

There are several example scenarios in which such multi-hop infrastructure access is useful and necessary. For example, consider an existing 802.11 access point deployed in a college campus as the infrastructure point. The access point could be installed near a building and connected to the wired Internet. Students may connect to the Internet while outside at some distance away from the building. In this scenario, they could use other students' laptops to reach the access point via multiple hops. Another example is that of personnel in a battlefield or rescue workers battling a disaster who may need to contact an infrastructure point, e.g., the field commander or the disaster control center. Another scenario is that of mobile robots teams that frequently need to communicate with the nest to send video and exploration data. Sensor networks also report sensor readings to a sink which is also a form of infrastructure. However, the use of aggregation in sensor networks differentiates such infrastructure access from the scenarios we are considering in this paper. The compressibility of sensor data improves the capacity in such networks despite the bottleneck at the sink [2] in contrast to the scenarios we investigate in this paper. Techniques for infrastructure access are also important in wireless mesh networks which aim to provide ubiquitous infrastructure access. Methods proposed in this paper can be used for mobile nodes to reach mesh routers that provide infrastructure access or for mesh routers to reach Internet gateways.

An important criterion for the success of infrastructure access is to ensure fairness among flows. However, there exist several challenges in providing fairness to nodes in in-

infrastructure access scenarios. First, transmissions of packets to and from the infrastructure point involves contention for channel access. Packets from one flow contend with those from other flows. Due to the many-to-one communication pattern, bottleneck links exist that limit the scheduling opportunities for non-contending flows. Second, the IEEE 802.11 MAC which is the accepted standard in wireless networking worsens the unfairness [13, 1]. The binary exponential backoff scheme favors the last node that successfully accessed the channel. This unfairness is illustrated by a simple scenario in which two nodes compete for channel access. The last successful node decreases its backoff interval while the unsuccessful node increases its backoff interval. There is a high probability of the first node succeeding in channel access repeatedly while the other node starves. Thus, some nodes tend to capture the channel while other nodes starve. Third, TCP tends to further worsen the unfairness problem in multi-hop wireless networks [4]. The retransmission timeout of starved nodes becomes larger and larger, resulting in a decrease in transport layer throughput. It has also been shown that there exists significant unfairness in flows crossing multi-hop wireless and wired networks [12]. In summary, due to the multi-hop nature of access to the infrastructure points, nodes that are further away have a low throughput due to the traversal of a larger number of hops and consequently unfair access. Note that this problem is fundamentally different from the extensively researched problem of providing fairness in one-hop access to infrastructure. In the one-hop case (e.g. WLANs), the infrastructure point can act as a coordinator and mitigate unfairness.

While ideally multi-hop fairness should be achieved among all the flows in the network, this is a hard problem to solve in a distributed fashion, especially when we consider that nodes may move continuously. In this paper, we propose that an important first step toward providing fairness in such multi-hop infrastructure access scenarios is to provide *symmetrical fairness*. Symmetrical fairness can be loosely defined as the problem of providing fair access to nodes that experience approximately equivalent network and channel conditions en route to the infrastructure point. This property indicates that although a flow with a larger hop count may be unfair in comparison to a flow with a smaller hop count, it should at least be fair with respect to other flows that are at similar distance from the access point and experience similar contention. Further, it can also be argued that a solution that provides symmetrical fairness is sufficient in multi-hop infrastructure access. This is because nodes that are farther away consume more network resources than nodes close to the infrastructure point, and thus if we consider the fairness in the usage of network resources, nodes that are close by should obtain higher throughput. In this context, providing fair access to nodes that consume sim-

ilar amount of resources when accessing the infrastructure is a more basic requirement. Further, we believe that symmetrical fairness is an important property for the usability of a multi-hop infrastructure access architecture. In the absence of the very basic (symmetrical) fairness, mobile users will tend to stop using the service (in commercial scenarios like multi-hop WLAN hotspots). In other applications where disconnection is not an option such as disaster relief, the absence of symmetrical fairness can hamper communication and coordination.

The contributions of this paper are threefold. First, we propose the notion of symmetrical fairness and a framework to characterize and measure it in multi-hop wireless networks. Second, we study the symmetrical fairness in infrastructure access in such networks and highlight several performance issues. Third, we propose and evaluate a distributed routing algorithm that provides symmetrical fairness in infrastructure access for multi-hop wireless networks.

2 Problem Description and Motivation

In this section, we first provide a framework to characterize symmetrical fairness. We then investigate a specific scenario of infrastructure access. We describe our model of this network and point out several performance issues observed.

2.1 Definition: Symmetrical Fairness

Our network model consists of a set of N mobile nodes each with a similar transmission range of R m and a bandwidth of W bps. Each mobile node has the same traffic rate and is continuously backlogged. Additionally, each mobile node originates traffic only to the infrastructure point, i.e. there is no peer-to-peer communication.

We define symmetrical fairness as the following property: nodes that are access-symmetric to each other should have similar throughput. Two nodes are access-symmetric if they experience similar levels of contention in accessing the infrastructure point, possibly through multiple hops. Thus the nodes in the network can be partitioned into a number of access-symmetric subsets AS_1 to AS_l . An AS (Access-Symmetric) subset is defined as a set of nodes whose throughput is expected to be similar. This is because all nodes that belong to a particular subset AS_i are defined to have similar access symmetry with relation to the infrastructure point and thus should ideally have similar throughput. The number of such subsets l ranges from 1 to N . For example, consider a network of 4 nodes surrounding a gateway. In this case there is only one AS subset $\{1, 2, 3, 4\}$. As another example, consider a chain network of 4 nodes

with the gateway placed at one end of the chain. In this scenario, there are 4 AS subsets and each node is in a unique AS subset.

Note that mobility will cause a node's membership in an AS subset to be transient. The node could change membership from one AS subset to another. Additionally, the number of subsets themselves could change. Thus with mobility it is hard to characterize the AS subsets over time. Our approach is to take a snapshot of the network and determine the AS subsets in the network at that time. The AS subsets determined in such a manner are assumed to exist for a duration of δt after the time when the snapshot is taken.

Once an AS subset is identified, we would like the throughputs of the nodes in that subset to be similar. We define a Subset Fairness Index (SFI) denoted by γ to characterize the fairness among the nodes in an AS subset. Let the throughput of node i be denoted as T_i . Then to identify the average fairness in each subset and its variance, we define

$$\gamma_{avg}(AS_i) = average\left(\frac{max(T_j, T_k)}{min(T_j, T_k)}\right), \forall j, k \in AS_i$$

$$\gamma_{var}(AS_i) = variance\left(\frac{max(T_j, T_k)}{min(T_j, T_k)}\right), \forall j, k \in AS_i$$

Similarly, we can characterize the worst case fairness in each subset as

$$\gamma_{max}(AS_i) = max\left(\frac{max(T_j, T_k)}{min(T_j, T_k)}\right), \forall j, k \in AS_i$$

Ideally, $\gamma(AS_i)$ for any AS_i should be 1. However due to symmetrical unfairness, it may be larger. In this manner we can define a symmetrical-fairness vector ρ for the whole network that contains l AS subsets as:

$$\rho = [\gamma_{avg}(AS_1), \gamma_{avg}(AS_2), \dots, \gamma_{avg}(AS_l)]$$

The ideal symmetrical fairness vector is denoted as ρ_{ideal} and for a network with l AS subsets is equal to:

$$\rho_{ideal} = [1, 1, \dots, 1]$$

Using these equations, we can define the overall fairness index FI of the network as:

$$FI = distance(\rho, \rho_{ideal})$$

where the distance function calculates the vector distance.

Thus, the problem of providing symmetrical fairness becomes minimizing FI. Note that it is a hard problem to characterize which nodes should be part of the same AS subset even when all nodes are assumed to have the same traffic rate. In this paper, we use a heuristic that looks at the hop

distance from the infrastructure point as well as contending neighbors at each hop to estimate which nodes should be part of the the same AS subset. For example, nodes that are the same hop distance away from the infrastructure point and have a similar number of contending neighbors can be part of the same AS subset. We use a top-down algorithm to determine the number and membership of AS subsets in the network. The top-down algorithm first constructs a shortest path tree rooted at the gateway that spans the wireless nodes. It then traverses down the tree in a breadth-first manner and partitions the nodes into subsets based on their hop counts and numbers of contending neighbors. Note that we do not use physical locations of nodes but rather a network-level metric like hop count which is representative of the performance achievable by a node. For example, packets from nodes that are longer hops away encounter higher delays due to queuing and contention at each hop.

Having defined how we characterize the symmetrical fairness, we now motivate the problem by showing that such symmetrical unfairness exists in infrastructure access for multi-hop wireless networks.

2.2 Symmetrical Fairness in Infrastructure Access

To motivate the symmetrical fairness problem, we investigate the performance of an example scenario for infrastructure access. In the rest of the paper, only as an example, we assume an infrastructure access scenario in which mobile nodes of a multi-hop wireless ad hoc network are trying to reach a node that has connectivity to the Internet. In this Internet access scenario, we refer to the infrastructure point as a gateway. In other words, we consider a scenario where mobile nodes use multi-hop routing to reach a gateway in order to access the Internet. We consider a mobile network of 49 nodes with one among them being the gateway. We use the modified random waypoint mobility model [15]. Nodes move with a speed uniformly distributed between 1 and 5 m/sec. The pause time is selected to be 0 second. Our simulation is performed in the Glomosim [16] simulator. The nodes are assumed to be placed uniformly in a grid fashion in an area of 1000mx1000m at $t = 0$ second and then subsequently move. The grid is a 7x7 grid with one node in each of them. Each node has a bit rate of 2Mbps and a radio range of 250m. The gateway is at the center of the area. Each node in the network initiates a UDP flow to the gateway with a packet rate of 3 packets per second. Each packet is 512 bytes in size.

Routes to the gateway are set up using a reactive approach in which a node that needs to access the gateway broadcasts a route discovery packet for the gateway, similarly as in the scheme described in [11]. However in our simulations, each node runs the DSR [8] routing algorithm.




Figure 1. Fairness in application scenario of infrastructure access in a multi-hop wireless network using the baseline algorithm.

Gateway discoveries are rebroadcast (ignoring duplicates) and the source route is embedded in the advertisement as it travels. The gateway on receipt of a route request sends a route reply to the requesting node. This node then stores the source route in the reply packet and caches it in a graph cache [7]. We assume each node has a means to realize that a destination is outside the wireless network (for example, based on address locality) and redirects such packets to the gateway. We will refer to this algorithm henceforth as Baseline. In order to measure the symmetrical fairness we use the algorithm discussed in Section 2.1 to split the nodes into subsets of nodes that are expected to be symmetrically fair. As the nodes are mobile, we take a snapshot of the network every 100 seconds and use the positions obtained to classify the nodes into AS subsets. Since all nodes generate equal amount of traffic and the node density is uniform, it is sufficient to use the hop distance in determining the AS subsets. We then measure γ_{avg} , γ_{var} and γ_{max} for each of the subsets. Figure 1 illustrates the symmetrical unfairness among nodes in this scenario. As can be seen from Figure 1(a), γ_{avg} is approximately 2. This indicates that nodes that should have symmetrical throughput are on an average 2 times less or more. Similarly, γ_{max} varies between 3 to 4 indicating that the lowest throughput of a node is 3 to 4 times smaller than the throughput of the other nodes. This shows evidence of severe symmetrical unfairness in the multi-hop Internet access scenario.

In the remainder of this paper, we investigate the causes of symmetrical unfairness in this network. We progressively propose solutions to overcome the causes and evaluate their impact on the performance. We begin our investigation by simplifying the Internet access scenario by assuming that the network setup above has no mobility.

3 Improving Infrastructure Access

In this section, we investigate the reasons for the lack of symmetrical fairness in the mobile multi-hop scenario setup in the previous section. In order to make the investigation

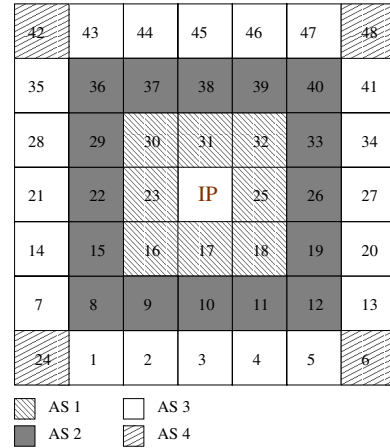


Figure 2. Access Symmetric subsets in infrastructure access.

more tractable, we assume that the nodes are no longer mobile. The remaining setup of the network is similar to in the previous section. Thus, in this section, we have a regular static network of 48 nodes and 1 gateway. Each node is placed in a grid and the whole network is split into 7x7 grids. The scenario is depicted in Figure 2. The AS subsets are shown individually shaded in Figure 2 and four such subsets exist in this scenario. The gateway in our experiments is placed in the center of the area. The placement of the gateway is orthogonal to the symmetrical unfairness problem. We have simulated two additional scenarios in which the gateway is placed at the corner and the top middle of the grid. In this scenario, the AS subsets need to be recalculated and their memberships are significantly different from the case in which the gateway is in the center of the network. However, the symmetrical fairness problem was found to exist regardless of the gateway placement.

We reran the Baseline algorithm in the static scenario. Figure 3 depicts the throughput of each node in the grid and Table 1 contains the fairness indices. The values of γ are similar to those in the mobile case as seen in Table 1. Thus, the static scenario experiences similar unfairness as in the mobile case.




Figure 3. Throughput performance with the Baseline algorithm. Severe unfairness and low throughput is observed.

AS Subset	AS1	AS2	AS3	AS4
Average	2.3	1.96	1.85	1.97
Variance	1.36	0.97	0.67	0.57
Worst Case	5	5.1	4.76	3.25

Table 1. Symmetrical fairness measure of the Baseline algorithm.

The symmetrical unfairness is further apparent from Figure 3. The first important observation is the high variance in throughput among different nodes. Nodes that are one hop away from the gateway have a slight imbalance in their throughputs while those further away have a large variation in their throughput. This is also indicated by the fairness indices of the subsets. In addition, nodes that are closer receive a higher share of the bandwidth than nodes further away. This is due to spatial bias and has been well documented in previous work. The second important observation is that there exists significant asymmetry in the throughputs observed. For example node in row 1 and row 7 are identically spaced away from the gateway and are likely to face similar amount of contention and queuing. However, they have noticeably asymmetric performance. This is captured by the fairness index too.

We now work toward improving the performance of the Baseline algorithm. In order to improve the performance of infrastructure access, we first investigate the reasons for the poor performance of the Baseline algorithm. In order to understand the dynamics of the network, we collected which routes were used by which node over time. One major reason for the poor performance observed is identified by this process. Since infrastructure access involves many-to-one communication, spatial reuse is very low since at some link along the route all flows in the network contend with each other. Thus, in such a network scenario a large

number of congestion/contention related packet drops occur. When such packet drops occur, the DSR protocol on the node that drops the packet sends a ROUTE ERROR back to the source of the packet. This creates more contention as well as causes the source nodes of the flow to remove the corresponding route to the gateway. Note that a valid route to the gateway is removed by this since there is no means to distinguish packets losses due to congestion/contention from losses due to wireless transmission errors. Now since the source node uses a graph based cache, it tries to reconstruct a new route to the gateway which is highly likely to have a higher hop count. Over time, each node tends to construct longer and longer routes (sometimes reaching 8 hops). Since the probability of packet error is higher for a path with a higher hop count, the performance gets worse over time. Additionally, since flows over longer hops are likely to be affected by packet drops more often, this problem affects the longer hop count flows more than the shorter ones.

In summary, despite reactive techniques or some hybrid of proactive and reactive techniques [11] having a lower overhead in terms of discovering routes, our experiments above have shown that the overall performance of the network can be adversely affected by the use of such techniques. One way to alleviate the problem due to reactive discovery is to modify DSR to disable the route cache and instead do a fresh route discovery when a route error is received. However, this approach can result in a large amount of overhead traffic. In the next section, we investigate whether a proactive modified gateway discovery technique can be used in such infrastructure access scenarios to alleviate the problem faced by reactive techniques.

3.1 Proactive Trusted Gateway Discovery

We propose to use the following proactive scheme for discovery and route maintenance to the gateway. The gateway in the network periodically broadcasts a GATEWAY ADVERTISEMENT packet. Each node in the network inserts its identifier (IP address) into a route record contained in the GATEWAY ADVERTISEMENT packet before rebroadcasting the GATEWAY ADVERTISEMENT packet. Each periodic GATEWAY ADVERTISEMENT contains a unique sequence number which is used by the nodes in the wireless network to refrain from rebroadcasting the GATEWAY ADVERTISEMENT more than once. Whenever a node receives a GATEWAY ADVERTISEMENT packet, it reverses the route record to form a source route pointing toward the gateway. This route is then added into the route cache. Subsequently, the freshest shortest route from this cache is used to deliver data packets to the gateway.

The frequency of the GATEWAY ADVERTISEMENT is a tradeoff between the staleness of routes and the overhead of




Figure 4. Throughput performance with proactive trusted gateway discovery.

AS Subset	AS1	AS2	AS3	AS4
Average	1.56	1.76	2.13	2.30
Variance	0.32	0.53	1.56	1.39
Worst Case	2.5	4.47	4.94	4.02

Table 2. Symmetrical fairness measure with proactive trusted gateway discovery.

advertisements. It should ideally be high enough to maintain valid routes in the presence of mobility. An important benefit of such advertisements is that nodes can assume that routes will remain valid until the next broadcast. Therefore, we propose to remove any network feedback (in the form of ROUTE ERRORS). Nodes *trust* the gateway to advertise routes periodically and thus can assume the route to be valid until the receipt of the next advertisement which will allow the node to populate the route cache with new routes to the gateway. We implemented the proactive trusted gateway discovery and simulated the performance of the network under this scenario. To make the investigation simple and decouple the effect of continuous broadcast on the overall congestion in the network, we allow the gateway to advertise for the first 10 seconds of the simulation and force the nodes to keep the routes obtained throughout the simulation. Since the network in this scenario is static, such an experiment is possible.

The results of this experiment are shown in Figure 4 and Table 2. The main observation is that the throughput of the network has increased significantly and the asymmetry is slightly reduced for subsets 1 and 2. However, significant asymmetry still exists. In order to investigate this, the routes used in the network were analyzed and are depicted in Figure 5. The center of each cell in the grid depicts the node at that location. The flows are depicted using arrows between the nodes. At the last hop before the gateway, the load (number of flows carried) of each one-hop flow from




Figure 5. Routes used by the various flows cause hotspots.

the one-hop neighbors of the gateway is shown. Note that there can be up to two outgoing arrows from a given node A . One arrow denotes the route advertised by A that may be used by nodes that choose A as a forwarding node. The other arrow denotes the route chosen by A from its route cache to forward its own packets. This route used by A depends on all the advertisements that arrived at A and may or may not be the same as the route advertised by A .

The main observation from Figure 5 is that the nature of routes selected cause hotspots in the network. Node 16 has one outgoing flow that has 13 multi-hop flows in it. Thus node 16 is not able to access the channel long enough for all the packets from these flows to get to the gateway. This results in the nodes using node 16 in their routes to have low throughput. This does not happen (by chance) in the upper part of the network due to which the throughput is more or less symmetric along each row and higher than that of the symmetrically opposite row.

3.2 Incorporating Route Diversity

In order to alleviate hotspots, we conjecture that incorporating some route diversity could be useful. In the previous section, only one set of routes were set up which would be used throughout. In this section, we allow the gateway to broadcast GATEWAY ADVERTISEMENT packets periodically. We denote the time period between two successive GATEWAY ADVERTISEMENT packets initiated by the gateway as the inter-advertisement period. This scheme results in a new flow graph being set up every time a GATEWAY ADVERTISEMENT packet is sent out. Thus, we expect that on average the performance should be better since hotspots may be formed in some inter-advertisement period while in other inter-advertisement periods, the network is likely to have symmetric routes. A node picks the shortest hop-count path found among multiple advertisements that it receives.

The results are shown in Figure 6 and Table 3. The results demonstrate that there is substantial improvement




Figure 6. Throughput performance with route diversity.

AS Subset	AS1	AS2	AS3	AS4
Average	1.48	1.56	1.49	1.65
Variance	0.31	0.24	0.18	0.21
Worst Case	2.56	2.51	2.39	2.47

Table 3. Symmetrical fairness measure with route diversity.

in the overall throughput and symmetrical fairness of the nodes as compared to the Baseline case. For instance, γ_{avg} for subset 1 improved from 2.3 in the Baseline case to 1.48 after incorporating route diversity. Similarly, all the other subsets showed improvement as can be seen in Table 3. However, asymmetry in the throughput performance still remains. This is because of two reasons. Firstly, for the nodes on the edges, the number of unique routes received from the gateway are fewer. Since they have fewer neighbors than a node closer to the center, they have fewer unique routes to the gateway. In addition, out of these smaller set of routes, an even smaller set of routes are shortest path routes. Thus, even if the GATEWAY ADVERTISEMENT packets arrive randomly, nodes may pick the same route resulting in a hotspot. Secondly, since routes are not timed out before the next GATEWAY ADVERTISEMENT is received, the links in the graph cache are not removed. Thus, when the node looks for a route, it may find an old route that was inserted in an earlier advertisement cycle. Therefore, multiple advertisements do not result in the use of new routes by nodes unless the old routes are timed out or removed. This problem is apparent in this scenario because the network is static. With mobility such problems are less likely to happen since old routes would be removed and new ones used.

In summary, having route diversity helps alleviate hotspots and makes the routes to be symmetrically set up for nodes that belong to the same subset. However, it does not solve this problem completely.




Figure 7. Flow graph of symmetric routes.




Figure 8. Throughput performance with symmetric routes.

3.3 Symmetric Routing

To ascertain whether symmetrically set up routes provide symmetrical fairness, we perform the following experiment in this section. We pre-configure symmetric routes in each node assuming global knowledge of the entire network. We define symmetric routes to be routes chosen for each node in the network such that nodes that are placed similarly with respect to the gateway use similarly loaded routes. The objective is to ascertain whether symmetrical unfairness still exists in such an idealized network scenario. The flow graph which is the union of all routes from all nodes to the gateway is depicted in Figure 7. The throughput performance of this scenario is depicted in Figure 8 and Table 4.

As seen from the results, the throughput of the nodes is fairly symmetrical in this idealized scenario. γ_{avg} for subset 1 and 4 is close to 1 and γ_{var} for these subsets is close to 0. This indicates close to ideal performance for these subsets. However, for subsets 2 and 3, the results are slightly higher than the ideal case because perfect symmetry in the load carried on each route cannot be pre-configured. In summary, symmetrical fairness can be ensured in the multi-hop scenario under the condition that nodes in the same subset pick

AS Subset	AS1	AS2	AS3	AS4
Average	1	1.35	1.38	1.08
Variance	0.0	0.18	0.13	0.0
Worst Case	1.01	1.99	1.86	1.14

Table 4. Symmetrical fairness measure with symmetric routes.

symmetric routes to the gateway. Furthermore, load balancing of these symmetric routes is also important to prevent hotspots. In the next section, we present a distributed algorithm to achieve this objective.

3.4 Random Routing Algorithm

Random Routing Algorithm (RRA) is a distributed algorithm designed to provide symmetrical fairness for infrastructure access among mobile nodes. RRA incorporates the findings from the previous studies in this paper and introduces a random routing component. We study two variants of this algorithm. Both RRA algorithms have the potential to provide symmetrical fairness in the infrastructure access scenarios that we describe. RRA is a completely distributed and low complexity protocol. It does not require the gateway to incorporate complex scheduling or tree building mechanisms to provide fairness. RRA is easily implementable due to its simplicity and involves only local decisions at each node.

The first algorithm, Source Routing RRA (SR-RRA), uses random selection of source routes. Each advertisement from the gateway results in each node receiving at least one unique route from each of its one hop neighbors. However, only some of these routes are minimum in hop count. Before sending each packet, SR-RRA identifies the set of minimum hop count routes from its cache of gateway routes. It then uniformly randomly chooses one of the routes from among these to send the current packet. In this manner each packet is sent by randomly selecting one of the minimum hop count routes available at the node. This technique has the following useful properties: (1) The use of random route selection reduces the probability of creating hotspots in certain areas of the network; (2) The routes used are independent of the order of their arrival. However, SR-RRA was found to still cause local hotspots. For example, node 3 in Figure 2 using SR-RRA may choose randomly between routes through nodes 9, 10 and 11. However all these routes could potentially go through node 17, resulting in a local hotspot. To alleviate such local hotspots, we propose a variant called Hop-by-Hop RRA (HbH-RRA). In HbH-RRA, each node chooses a random next hop node from among its set of minimum hop count routes to the gateway. Thus, instead of specifying source routes, each nodes makes a local random choice at every hop to forward a packet toward a gateway. In SR-RRA, randomly selecting a route only




Figure 9. Throughput performance with the Random Routing Algorithm (HbH-RRA).

AS Subset	AS1	AS2	AS3	AS4
Average	1.35	1.17	1.13	1.1
Variance	0.08	0.03	0.01	0.0
Worst Case	1.69	1.52	1.38	1.2


Table 5. Symmetrical fairness measure with HbH-RRA.

allows for random selection of the first hop toward the gateway. However, all routes stored may still go through the same second hop node creating a local hotspot. In contrast to SR-RRA, nodes 9, 10 and 11 in HbH-RRA can randomly select a next hop that may or may not be node 17, alleviating the local hotspot at node 17. Thus HbH-RRA reduces the probability of local hotspots by allowing each node to forward to a random next hop that makes progress toward the gateway.


Figure 9 and Table 5 depict the results of HbH-RRA. As can be seen, HbH-RRA further brings γ_{avg} closer to the ideal value of 1 for each subset. For instance, for subset 4, γ_{avg} is reduced from 1.65 to 1.1. Similarly, γ_{var} is also reduced and is close to the ideal value of 0. Thus, RRA successfully improves the symmetrical fairness obtained by each node.

3.5 Summary of Improvements


This section summarizes the progressive improvement obtained from the techniques employed to improve symmetrical fairness. Figure 10 shows the mean SFI (Figure 10(a)) and variance (Figure 10(b)) of SFI for the different techniques over different subsets. We started with a regular static network using our Baseline algorithm (nodes running DSR contacting the gateway) and observed the presence of symmetrical unfairness as defined by our framework. The first problem solved was of removing false feed-



(a) Mean SFI



(b) Variance of SFI



(c) Throughput

Figure 10. Summary of improvements in symmetrical fairness.

back (inferring congestion as link failures). This improves the mean and variance of SFI compared to the Baseline algorithm significantly. The second problem solved was allowing for route diversity which showed a modest gain in symmetrical fairness. Finally, the RRA algorithm shows a symmetrical fairness competitive with ideal symmetrical routing. RRA provides a good fairness measure close to that of symmetric routing which assumes global knowledge. Although both RRA algorithms provide improved symmetrical fairness, HbH-RRA provides a lower mean and variance of SFI across all AS subsets as compared to SR-RRA. Another important issue to consider is whether throughput (aggregate and AS-level) suffers while improving symmetrical fairness. Throughput is defined as the percentage fraction of packets received at the gateway from among those transmitted by a node. Figure 10(c) shows that HbH-RRA always improves or maintains similar throughput in comparison to the other schemes. In fact, the overall throughput is improved from 42% to 71% using HbH-RRA while the AS-level throughputs are also improved as compared to the Baseline algorithm. In conclusion, HbH-RRA provides the best symmetrical-fairness while improving the throughput of the Baseline algorithm in infrastructure access.

3.6 Impact of Mobility

In this section, we evaluate the performance of HbH-RRA in the presence of mobility in multi-hop wireless infrastructure access by repeating the same experiment in Section 2.2 but using HbH-RRA. As seen from Table 6, the symmetrical fairness of nodes is improved when using HbH-RRA as compared to using the Baseline algorithm. The improvement is evident from the reduction in the mean and the variance of the SFI. Thus, HbH-RRA successfully improves the symmetrical fairness even in a mobile scenario. We also found that the throughput improved by 5% using HbH-RRA as compared to using the Baseline algorithm.

4 Related Work

Several works [10, 6, 9] tackle the fairness problem in multi-hop wireless networks. However, the solutions proposed involve modifications to the MAC layer. Our focus in this work is to provide symmetrical fairness with a simple and readily implementable routing layer solution that can be implemented on off-the-shelf hardware and requires no modification to firmware.

The work in [14] is related to one of the scenarios we investigate in this paper. The paper proposes a solution to improve the fairness among competing TCP flows crossing wireless and wired networks. The scenario the authors consider is similar to a scenario for infrastructure access to reach the Internet. Their solution involves delaying packets

	AS1		AS2		AS3		AS4	
	Baseline	HbH-RRA	Baseline	HbH-RRA	Baseline	HbH-RRA	Baseline	HbH-RRA
Average	2.07	1.35	1.49	1.25	1.65	1.15	1.91	1.11
Variance	1.78	0.34	0.47	0.24	0.62	0.19	1.01	0.14
Maximum	4.08	2.56	3.39	2.21	4.13	2.10	3.02	2.01

Table 6. Fairness in application scenario of infrastructure access in a multi-hop wireless network, using Baseline and HbH-RRA algorithms.

of nodes that observe their own output to be high. In [3], the authors propose solutions toward providing multi-hop fairness (longer hop nodes have similar throughput to shorter hop nodes) in static mesh networks. However, the operation of their algorithm in a mobile and dynamic topology is not evaluated and the topology is limited to only a single chain of nodes with known traffic patterns. Additionally, fairness algorithms for mesh networks can have different assumptions and tradeoffs since mesh routers are similar to infrastructure devices. The work in [5] provides algorithms for dynamic load balancing in wireless access networks similar to mesh networks. However, they assume independent interfaces to every neighbor which removes any contention during medium access. Additionally, only end nodes are assumed to be mobile.

5 Conclusions

In this paper, we propose a symmetrical fairness model for infrastructure access in multi-hop wireless networks. Symmetrical fairness is an important step toward providing fairness in such multi-hop wireless networks. We provide a framework for characterizing and measuring the symmetrical fairness in a given network and define fairness indices to compare different routing schemes.

Using detailed simulations we show the existence of symmetrical unfairness in a typical infrastructure access scenario. We then provide solutions for improving the symmetrical-fairness for infrastructure access. Finally, we present a distributed routing algorithm called RRA that improves symmetrical fairness in such infrastructure access scenarios in mobile multi-hop wireless networks. Our results demonstrate that RRA improves symmetrical-fairness without any loss in the overall throughput or AS-level throughput when compared to the Baseline algorithm.

Our future work involves quantifying the benefit of multiple infrastructure points in the network as well as schemes toward load balancing and efficient gateway selection in such scenarios. Another important issue is to incorporate multi-hop fairness among flows by removing spatial bias, i.e., the flows with longer hop counts should have similar throughput to the flows with lower hop counts. We are particularly interested in the performance of such fairness algorithms in the presence of mobility.

Acknowledgment

This work was supported by NSF grant ANI-0338856.

References

- [1] V. Bhargavan, A. Demers, S. Shenker, and L. Zhang. Macaw: A media access protocol for wireless lans. *Proc. of ACM SIGCOMM*, 1994.
- [2] E. J. Duarte-Melo and M. Liu. Data-gathering wireless sensor networks: organization and capacity. *Comput. Networks*, 43(4):519–537, 2003.
- [3] V. Gambiroza, B. Sadeghi, and E. W. Knightly. End-to-end performance and fairness in multihop wireless backhaul networks. In *Proc. of ACM MobiCom*, 2004.
- [4] M. Gerla, K. Tang, and R. Bagrodia. Tcp performance in wireless multi-hop networks. In *Proc. of WMCSA*, 1999.
- [5] P. H. Hsiao, A. Hwang, H. T. Kung, and D. Vlah. Load-balancing routing for wireless access networks. In *Proc. of IEEE INFOCOM*, 2001.
- [6] H.-Y. Hsieh and R. Sivakumar. Improving fairness and throughput in multi-hop wireless networks. In *Proc. of ICN*, July 2001.
- [7] Y.-C. Hu and D. B. Johnson. Caching Strategies in On-Demand Routing Protocols for Wireless Ad Hoc Networks. In *Proc. of ACM MobiCom*, August 2000.
- [8] D. B. Johnson and D. A. Maltz. *Dynamic Source Routing in Ad Hoc Wireless Networks*. Kluwer Academic, 1996.
- [9] H. Luo, S. Lu, V. Bhargavan, J. Cheng, and G. Zhong. A packet scheduling approach to qos support in multihop wireless networks. *Mob. Netw. Appl.*, 9(3):193–206, 2004.
- [10] T. Nandagopal, T.-E. Kim, X. Gao, and V. Bhargavan. Achieving mac layer fairness in wireless packet networks. In *Proc. of ACM MobiCom*, 2000.
- [11] P. Ratanchandani and R. Kravets. A hybrid approach for internet connectivity for mobile ad hoc networks. In *Proc. of IEEE WCNC*, 2003.
- [12] K. Xu, S. Bae, S. Lee, and M. Gerla. Tcp behavior across multihop wireless networks and the wired internet. In *Proc. of WoWMoM*, 2002.
- [13] S. Xu and T. Saadawi. Does the IEEE 802.11 mac protocol work well in multihop wireless ad hoc networks? *IEEE Communications Magazine*, 39(6):130–137, 2001.
- [14] L. Yang, W. K. Seah, and Q. Yin. Improving fairness among tcp flows crossing wireless ad hoc and wired networks. In *Proc. of ACM MobiHoc*, 2003.
- [15] J. Yoon, M. Liu, and B. Noble. Random waypoint considered harmful. In *Proc. of IEEE INFOCOM*, April 2003.
- [16] X. Zeng, R. Bagrodia, and M. Gerla. Glomosim: A library for parallel simulation of large-scale wireless networks. In *Proc. of PADS Workshop*, May 1998.