CHAPTER 6

SUPPORTING HIGH-THROUGHPUT ROUTING METRICS

6.1 Introduction

Researchers have proposed a number of metrics to select routes with high throughput in multi-hop wireless networks. De Couto et al. [31] proposed the ETX (expected transmission count) metric. The ETX of a link is the expected number of transmissions to send a packet over that link, and the ETX of a path is the sum of the ETX of each link on the path. Draves et al. [34] proposed the ETT (expected transmission time) metric to find high throughput paths in heterogeneous, multiradio WMNs. In this chapter, we collectively call metrics that select routes with high throughput in multi-hop wireless networks High-Throughput Metrics (HTMs). Researchers prefer on-demand routing protocols in wireless ad-hoc networks, because they save resources such as bandwidth and processing power [23, 57]. To use HTMs in on-demand routing protocols, typically metrics of all links on a path need to be collected [31, 34]. Thus, protocols that use source routing are suitable to support HTMs, since they contain individual link information in their route-discovery messages. As a result, DSR (dynamic source routing) [58] is a natural choice when supporting an HTM. Both ETX and ETT have been implemented as a metric in DSR [31, 34]. However, the existing design of supporting HTMs in DSR is not efficient, because it collects link state information and runs Dijkstra's shortest path algorithm.

AODV (ad hoc on-demand distance vector) [85] is another well-known and well-studied ondemand routing protocol. However, when trying to support an HTM on AODV, we face problems because control messages in AODV typically do not contain individual link information. ETX is also implemented as a metric in

DSDV (destination-sequenced distance-vector routing [84]) in [31]. However, DSDV works proactively instead of reactively. AODV-ST (spanning tree) [90] supports the ETT metric, but it is based on infrastructure wireless networks. Therefore, the motivation of this work is to design an efficient and generalized scheme to support HTMs in on-demand routing protocols, including AODV [120].

6.2 High-Throughput Metrics (HTMs)

A number of HTMs to measure the qualities of wireless links and wireless paths composed of multiple hops. We can classify these metrics according to different criteria, reflecting their fundamental design and implementation choices [121].

6.3 The Existing Approach to Supporting HTMs in DSR

DSR (dynamic source routing) [58] is a natural choice in supporting HTMs in on-demand routing protocols (e.g., [31, 34, 33]), because an HTM scheme typically needs to collect link state information to calculate the shortest path. In DSR, a source node explicitly designates the route to a destination in the header of data packets. Routes are established by recording traversed paths in RREQ (route request) messages and returning discovered paths in RREP (route reply) messages [122].

6.4 APM: Supporting HTMs in On-Demand Routing Protocols

Routes discovered by APM are the shortest paths. When supporting an HTM, AODV has difficulties because its route discovery messages do not contain individual link information. Thus, it is not able to collect link state information and run Dijkstra's shortest path algorithm. A possible solution is to let RREQ and RREP messages in AODV carry the entire path. However, this approach is not desirable because it deprives AODV of its simplicity. For example, RREQ and RREP messages would no longer be small and constant in size [123].

From the route discovery process of DSR-HTM, we make two observations:

First, intermediate nodes propagate not only the first received RREQ packet, but also duplicate RREQ packets that have a new better (shorter) path.

Second, to determine if a path is shorter than another, only the accumulated path metric, instead of individual link metrics, is used for comparison.

6.5 Path Throughputs

Figure 6.1 compares the throughput of routes found with a minimum hop-count metric to the throughput of the best routes that could be found. Each curve shows the throughput CDF (in packets per second) for 100 node pairs; the pairs are randomly selected from the 29×28=812 total ordered pairs in the test-bed. A point's x value indicates throughput, in packets per second; they value indicates what fraction of pairs had less throughput. The left curve is the throughput CDF achieved by routing data using DSDV with the minimum hop-count metric. The right curve is the throughput CDF for the best known route between each pair of nodes. Packets were only sent between one pair at a time. [124] For each pair, the DSDV and best-path tests were run immediately after one another, to limit variation in link conditions over time. The "best" path between each pair of nodes was found by sending data along ten potential best paths, one at a time, and selecting the path with the highest throughput. Potential best paths were identified by running an off-line routing algorithm, using as input measurements of per-link loss ratios and with a penalty to reflect the reduction in throughput caused by interference between successive hops of multi-hop paths. New link measurements were collected roughly every hour during the experiment; the best paths for each pair were generated using the most recently available loss data. The values in Figure 6.1 are split into two main ranges [125], above and below 225 packets per second. The values above 225 correspond to pairs that communicated along single-hop paths; those at or below 225 correspond to multi-hop paths. A single-hop direct route can deliver up to about 450 packets per second, but the fastest two-hop route has only half that capacity. The halving is due to transmissions on the successive hops interfering with each other: the middle node cannot receive a packet from the first node at the same time it is sending a packet to the final node [126]. Similar effects cause the fastest three-hop route to have a capacity of about 450/3 = 150 packets per second.

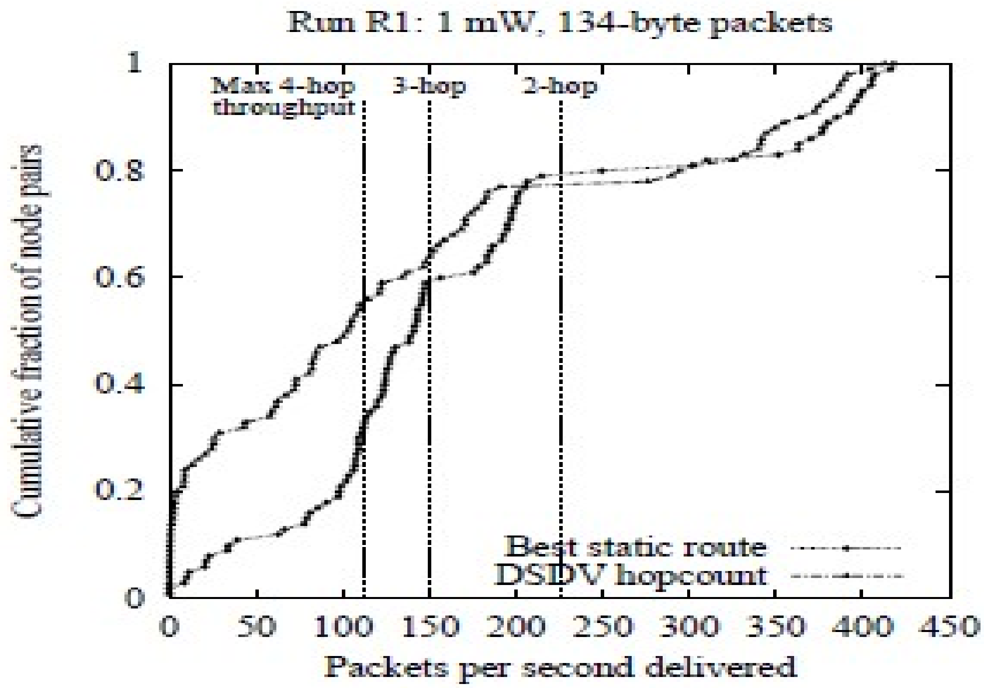

Figure 6.1: When using the minimum hop-count metric, DSDV chooses paths with far less throughput than the best available routes. Each line is a throughput CDF for the same 100 randomly selected node pairs. The left curve is the throughput CDF of DSDV with minimum hop-count. The right curve is the CDF of the best throughput between each pair, found by trying a number of promising paths. The dotted vertical lines mark the theoretical maximum throughput of routes of each hop-count.

Minimum hop-count performs well whenever the shortest route is also the fastest route, especially when there is a one-hop link with a low loss ratio. A one-hop link with a loss ratio of less than 50% will outperform any other route [127]. This is the case for all the points in the right half of Figure 2. Note that the overhead of DSDV route advertisements reduces the maximum link capacity by about 15 to 25 packets per second, which is clearly visible in this part of the graph. The left half of the graph shows what happens when minimum hop-count has a choice among a number of multi-hop routes. In these cases, the hop-count metric usually picks a route significantly slower than the best known [128]. The most extreme cases are the points at the far left, in which minimum hop-count is getting a throughput close to zero, and the best known route has a throughput of about 100 packets per second. The minimum hop-count routes are slow because they include links with high loss ratios, which cause bandwidth to be consumed by retransmissions.

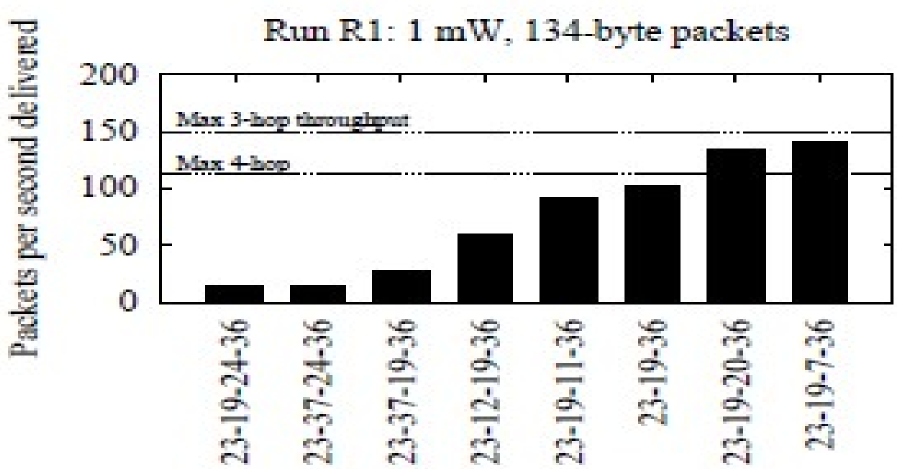

Figure 6.2: Throughput available between one pair of nodes, 23 and 36, along the best eight routes tested. The shortest of the routes does not perform the best, and there are a number of routes with the same number of hops that provide very different throughput.

6.6 Distribution of Path Throughputs

Figure 6.2 illustrates a typical case in which minimum hop-count routing would not favor the highest-throughput route. The through-put of eight routes from node 23 to node 36 is shown [129]. The routes are the eight best which were tested in the experiments described above. The graph shows that the shortest path, a two-hop route through node 19, does not yield the highest throughput. The best route is three hops long, but there are a number of available three-hop routes which provide widely varying performance. A routing protocol that selects randomly from the shortest hop- count routes is unlikely to make the best choice, particularly as the network grows and the number of possible paths between a given pair increases [130].

6.7 Summary

In this chapter we will discuss the supporting high throughput routing metrics. Researchers have proposed a number of metrics to select routes with high throughput in multi-hop wireless networks.

In This Chapter we will discuss the following points are:

- High-Throughput Metrics (HTMs)
- The Existing Approach to Supporting HTMs in DSR
- APM: Supporting HTMs in On-Demand Routing Protocols
- Path Throughputs, Distribution of Path Throughputs