Available online at www.sciencedirect.com
SCIENCE DIRECT®
Ad Hoc Networks 4 (2006) 303–325
Ad Hoc Networks
www.elsevier.com/locate/adhoc

Routing characteristics of ad hoc networks with unidirectional links

Jorjeta G. Jetcheva [a,*], David B. Johnson [b]

[a] Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213, United States
[b] Computer Science Department, Rice University, Houston, TX 77005, United States

Received 11 March 2004; accepted 11 June 2004
Available online 21 October 2004

Abstract

Unidirectional links in an ad hoc network can result from factors such as heterogeneity of receiver and transmitter hardware, power control or topology control algorithms, or differing sources of interference or jamming. Previously proposed metrics for evaluating the difficulty of a unidirectional scenario are limited in scope and are frequently misleading. To be able to analyze ad hoc network routing protocol behavior in a complex networking environment, it is not sufficient to merely assign a single level of difficulty to a unidirectional network scenario; the many interrelated *routing characteristics* of these scenarios must be understood. In this paper, we develop a set of metrics for describing these characteristics, for example for characterizing routing scenarios in simulations, analysis, and testbed implementations. Based on these metrics, we perform a detailed simulation analysis of the routing characteristics of the three most common simulation models for generating unidirectional links in ad hoc networks: the *random-power model*, the *two-power model*, and the *three-power model*. Our findings enable protocol designers to better choose a set of network scenarios and parameters that truly explore a wide range of a routing protocol's behaviors in the presence of unidirectional links, and to better understand the complex interplay between routing mechanisms and network conditions.

Keywords: Unidirectional links; Routing; Ad hoc networks; Mesh networks

1. Introduction

A group of wireless mobile hosts that wish to communicate may form an *ad hoc network*, forwarding packets for each other to support hosts beyond the single-hop wireless transmission range. Most routing research in wireless ad hoc networks

* Corresponding author. Tel.: +1 408 399 7771.
E-mail addresses: jorjeta@cs.cmu.edu (J.G. Jetcheva), dbj@cs.rice.edu (D.B. Johnson).

1570-8705/$ - see front matter
doi:10.1016/j.adhoc.2004.09.001

has been based on the simplifying assumption that all wireless links in the network are bidirectional and would therefore work equally well in both directions. However, there exist a variety of circumstances in which this assumption does not hold. *Unidirectional* links can result from factors such as heterogeneity of receiver and transmitter hardware (leading to differing transmission ranges), power control algorithms (in which nodes vary their transmission power based on their current energy reserves), or topology control algorithms (aimed at reducing interference in the network by computing the lowest transmit power that each node needs to stay connected to the network). Unidirectional links may also result from interference around a node that prevents it from receiving packets even though other nodes are able to receive its packets. For example, unidirectional links may often occur if some nodes in the ad hoc network are mounted on vehicles, using powerful transmitters, and others are carried by pedestrians, using battery-powered transmitters with a shorter transmission range.

Routing protocols for ad hoc networks may deal with the possibility of unidirectional links in a number of ways:

- Some protocols, such as DSDV [9] or conventional distance vector routing protocols, simply do not consider the problem and thus may create routes that fail to work, causing packet losses.
- Other protocols handle the presence of unidirectional links by treating all links as if they *might* be unidirectional, in order to avoid the problem above, but this causes extra routing overhead. For example, DSR [5,6] normally returns a Route Reply packet using the reverse of the route recorded in the Route Request packet, but this only works over bidirectional links; DSR alternatively can be configured to independently discover the route for returning the Route Reply, which can significantly increase routing overhead.
- Still other protocols detect and keep track of unidirectional links as network conditions change, and then either use them for routing (e.g., [11]) or simply ignore them as in AODV [10]. Nodes using AODV attempt to learn which links to neighboring nodes are unidirectional; such a neighbor is remembered in a "blacklist" set, and new routes through such neighbors are ignored and not used for some time. This approach limits connectivity in the network, and may lead to increased overhead and failure to setup a route, when the mechanism to learn of unidirectional links has not yet detected a new unidirectional link or when bidirectional communication becomes possible over a previously unidirectional link.

There is growing interest in routing techniques that enable efficient routing over networks with unidirectional links, and in the effects of unidirectional links on routing protocol performance (e.g., [1,7,8,11,14]). However, little work has been done on evaluating the routing difficulty or *routing characteristics* of ad hoc networks with unidirectional links, for instance for characterizing routing scenarios in simulations, analysis, and testbed implementations. In general, routing difficulty is higher when the likelihood that the routing protocol would encounter a unidirectional link is higher, since handling each unidirectional link requires additional mechanisms and control packet transmissions to either ignore or be able to use the link for routing. For example, if some node **A** receives a packet directly from a node **B** but the link from **B** to **A** is unidirectional, then **A** would not be able to return an acknowledgment packet directly to **B**; rather, node **A** would need to send the acknowledgment to **B** along a multihop path (a multihop acknowledgment).

In prior work, the difficulty that a unidirectional scenario presents to the routing protocol has frequently been expressed simply as the number or fraction of unidirectional links in the network. However, these metrics present a very limited view of the routing characteristics of the network and do not adequately reflect the routing difficulty of a unidirectional network scenario. For example, AODV [10] is an on-demand routing protocol for ad hoc networks that utilizes Route Request packet floods to discover routes and to perform route repair when routes break. Fig. 1 shows the number of Route Request transmissions in AODV for

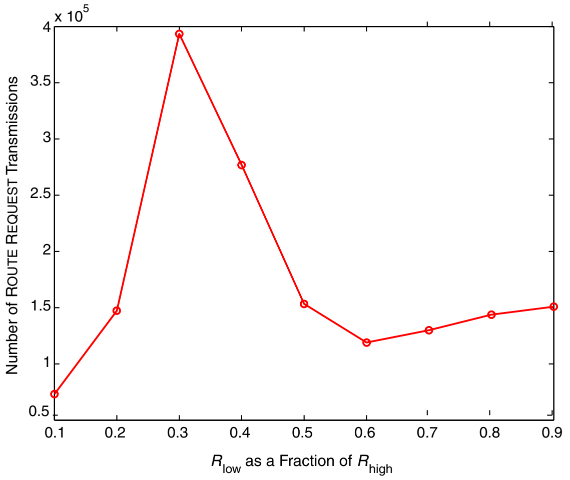

Fig. 1. Simulation of AODV in mobile ad hoc networks with 10 nodes with transmission range $R_{high} = 250$ m and 90 nodes with range R_{low}, all moving according to a Random Waypoint model with 1 m/s maximum speed and 0 pause time. 35 pairs of nodes exchange one-way CBR traffic at four 64-byte packets per second.

a set of unidirectional scenarios. Since the number of unidirectional links and the fraction of unidirectional links decreases monotonically from left to right in the graph (in-neighbors curve in Fig. 2), the effort expended by the routing protocol, as expressed by the number of Route Request transmissions, might be expected to also decrease monotonically from left to right. However, this is not the case. We are able to analyse the behavior of AODV in Section 5.7 with the help of the metrics we develop in Section 3.

In this paper, we develop a set of metrics to expose the routing characteristics of networks with unidirectional links. Although some of these metrics have been considered separately in prior work, the combined set of metrics is the first to provide a comprehensive view of unidirectional networks from the point of view of routing protocol design and evaluation.

We demonstrate the usefulness of our set of metrics by performing detailed simulation-based evaluation of the routing characteristics of networks with unidirectional links generated by the three most common simulation models for generating unidirectional links in ad hoc networks: the *random-power model*, the *two-power model*, and the *three-power model*. We use a wide range of parameterizations of each model (e.g., power levels and number of nodes per power level) and generate both mobile and static unidirectional network scenarios. In the course of our evaluation, we show how commonly used sets of metrics may be misleading and may fail to reveal some important characteristics of the network. We base this evaluation on simulation rather than on formal analysis, in order to account for the detailed, complex behavior of the network and individual packets, without the abstractions and simplifications often necessary with formal analysis.

Our metrics and analysis of the characteristics of the three power models enable protocol designers to better choose a set of network scenarios and parameters that truly explore a wide range of a routing protocol's behaviors in the presence of unidirectional links, and to better understand the subtleties of the interplay between routing mechanisms and network conditions.

The rest of this paper is organized as follows. Section 2 discusses related work. We present the power variation simulation models in Section 3. In Section 4, we present our routing characteristics metrics, and in Section 5, we present our simulation results. Section 6 concludes the paper.

2. Related work

Most prior work on routing in the presence of unidirectional links uses unidirectional network scenarios generated by a random-, two-, or three-power model. The difficulty or routing characteristics of these scenarios are often not evaluated or are evaluated using only one or two metrics. In addition, some metrics used in previous work, although useful in some ways in characterizing the properties of an ad hoc network, were not specifically designed to expose its routing characteristics.

Pearlman et al. [8] presented mechanisms for routing over unidirectional links and evaluated these mechanisms in networks generated by a two-power model. However, they analyzed the routing difficulty of these network scenarios by looking only at the fraction of unidirectional links in each. We build on their work by studying a larger set of parameterizations of the two-power model, and by conducting a comprehensive evaluation of the routing characteristics of the scenarios generated.

Sinha et al. [14] proposed extensions to the Zone Routing Protocol (ZRP) [4] for handling unidirectional links and evaluated these extensions in networks with unidirectional links produced using the same two-power approach but with less parameterizations and with no characterization of the unidirectional scenarios.

Ramasubramanian et al. [11] presented a protocol for a bidirectional abstraction of unidirectional links that handles them below the routing layer. The authors used five unidirectional network scenarios to evaluate the proposed protocol, generated using a two-power, a three-power, and a four-power model. The parameters for generating these scenarios were hand-picked to provide a range of routing difficulty. The authors use only the average number of in-neighbors of a node (neighbors that can reach a node but cannot be reached by it) and average reverse route length to characterize these scenarios.

Marina and Das [7] investigated the potential benefits of using unidirectional links for routing in ad hoc networks, using a random-power model, a two-power model, and a distributed topology control method. They characterized the scenarios generated by each model by analyzing the connectivity of the networks in each scenario in terms of the size and number of strongly connected components in networks with different densities. They also measured the difference in path length between a network scenario with unidirectional links and the same scenario with only bidirectional links. With the exception of the path metric, the rest of the metrics were not specifically designed to study the routing characteristics of the network. Our analysis is based on a wider range of parameterizations of the random- and two-power models, uses a set of metrics to expose the routing characteristics of the network, and also includes the three-power model, not discussed by Marina and Das [7]. Unlike their work, we do not discuss topology control algorithms here, as each topology control algorithm is unique, and it is thus difficult to generalize the results from such analysis.

Ramasubramanian and Mosse [12] also studied the connectivity of ad hoc networks with unidirectional links. Their analysis focused only on static networks with power variations produced by a random-power model and by a model that determines the directionality of a link according to a uniform random distribution (e.g., each link has a probability P of being unidirectional). The metrics they used to characterize network connectivity are reverse route length, the average size and distribution of connected components for reverse route lengths smaller than 1, 2, 3, or 4 hops, and average path length. In this paper, we use a broad range of metrics to expose the characteristics of ad hoc networks with unidirectional links from a routing perspective in the context of three different power models.

3. Methodology

In this section, we describe the network parameters and power models that we use to generate ad hoc networks with unidirectional links.

3.1. Network parameters

We simulated networks of 100 mobile nodes using *ns-2* [3]. The nodes were placed on a rectangular 1200 m $\times$ 800 m area. A rectangular area was chosen in order to allow for the formation of longer paths (along the longer dimension) than would be possible with a square site having the same total area.

The nodes move according to the Random Waypoint model [2]: each node is initially placed at a uniformly randomly chosen location, stays there for a period of time called the *pause time*, and then starts moving toward a new randomly chosen location, at a randomly chosen speed up to some maximum speed; once the node reaches this new location, it remains stationary again for the pause time and then repeats the process, moving toward a new location at a new speed.

Each simulation was run for 900 s of simulated time. We use two maximum speeds, 1 m/s and 20 m/s, and two pause times, 0 s (continuous mobility) and 900 s (a static network).

Wireless transmission attenuation follows a free-space model up to a reference distance and then a two-ray ground reflection model [2]. The nominal (omni-directional) range of each device is 250 m.

These network parameters are representative of those widely used in performance analysis of ad hoc networks.

The simulation environment makes the simplifying assumption that all nodes have a circular wireless transmission pattern. We believe that our results are also useful in describing the effect of unidirectional links on routing protocol behavior with more generalized transmission patterns that also produce unidirectional links.

3.2. Power variation models

The most commonly used methods for generating unidirectional network scenarios for use in ad hoc network simulations are based on varying the transmission range of nodes in the network, whereby the transmission range of each node (or set of nodes) is set to a fraction of the nominal transmission range.

Two characteristics of these models make them attractive for use in simulations of ad hoc networks. First, the resulting scenarios are realistic, as they correspond to a potentially common situation in which the ad hoc network consists of heterogeneous devices that may have different transmitter and/or receiver capabilities. Second, the scenarios are straightforward to generate.

For simplicity, in the rest of this paper, wireless range is discussed in terms of distance instead of transmission power level, because distance is easier to conceptualize.

3.2.1. Random-power model

In the random-power model, each node is assigned a random transmission power level chosen from a range of transmission powers that correspond to a (R_{min}, R_{max}) distance range. In our work, we use R_{max} of 250 m, the nominal transmission range, and we study values of R_{min} ranging from 0 m to 225 m at 25 m increments. This model generates networks that are "completely" heterogeneous, in that each device may have a different transmission range.

3.2.2. Two-power model

In the two-power model, the nodes in the network are divided into two groups, each of which is assigned a different power level, corresponding to a different transmission range, R_{low} and R_{high}, respectively; N_{low} total nodes have range R_{low}, and N_{high} total nodes have range R_{high} (set to the nominal transmission range). We vary R_{low} from 0 m to 225 m at 25 m increments, and vary N_{low} from 10% to 90% of the total number of nodes.

3.2.3. Three-power model

In the three-power model, the nodes in the network are divided into three groups, each of which is assigned a different power level, corresponding to a different transmission range, R_{low}, R_{medium}, and R_{high}, respectively. N_{low} total nodes have range R_{low}, N_{medium} total nodes have range R_{medium}, and N_{high} total nodes have range R_{high} (set to the nominal transmission range). In this evaluation, N_{low} and N_{medium} were varied between 10% and 90% of the total number of nodes, at 10% increments, thus covering the range of possible combinations of (N_{low},N_{medium},N_{high}) at a 10% discretization. The simulation computational cost of varying R_{low} and R_{medium} to cover all possible combinations even at a 10% discretization is prohibitive. Instead, two sets of values for R_{low} and R_{medium} are used in our experiments, and only the fraction of nodes in each power group was varied. The values for R_{low} and R_{medium} are 0.6 and 0.7 of the nominal range for the first set of network scenarios (which we call Scenario 1), and 0.2 and 0.4 for the second set of network scenarios (which we call Scenario 2). The differences between the values of any pair in the set R_{low}, R_{medium}, R_{high} in Scenario 1 are twice as large as the corresponding differences in Scenario 1. Our experiments do not cover a full set of power parameterizations of the three-power model, but Scenarios 1 and 2 explore two different points in its behavior.

4. Metrics

In this section, we propose a set of metrics to expose the routing characteristics of unidirectional network scenarios. These metrics cover a wide range of network events and states as expressed by the primitives that directly affect routing protocol performance: links and routes (or paths).

Since most current routing protocols strive to route along the shortest path between a source and a destination, all path-related metrics in the analysis consider only the shortest paths between each pair of nodes that exist in the network, as these paths are the most likely to be used by the routing protocol.

All metrics are computed as averages over all nodes in the network and over the lifetime of the simulation.

4.1. Neighbor-related metrics

The number of *bi-neighbors* of a node is the number of neighboring nodes (i.e., nodes within transmission range of the node) with which a node has bidirectional links. The *in-neighbors* of a node are the neighbors of the node that can reach the node but cannot be reached by it. The sum of the in- and bi-neighbors of a node is the total number of neighbors of the node.

The *total number of neighbors* metric reflects the level of connectivity in the network. For example, the higher the total number of neighbors per node, the higher the number of paths along which nodes can communicate with each other (on average).

The *in-neighbors* metric reflects the average number of unidirectional links that are present in the network.

The *in-neighbors fraction* metric represents the number of in-neighbors as a fraction of the total number of neighbors. It indicates the likelihood that the routing protocol will encounter a unidirectional link rather than a bidirectional link. Each encountered unidirectional link may cause the routing protocol additional overhead to either use the link or to avoid it, as described in Section 1.

4.2. Node reachability metrics

The *unreachable nodes* metric represents the average number of nodes to which a node does not have a route and which do not have a route to this node (no routes exist). This metric reflects the extent to which the network is partitioned and can help protocol designers to differentiate between poor routing protocol performance due to the presence of unidirectional links versus poor routing protocol performance due to the presence of partitions.

The *bidirectionally reachable nodes* metric represents the number of nodes to which the average network node has a shortest path that is bidirectional. This metric reflects the likelihood that the routing protocol would encounter a bidirectional path for routing between a pair of nodes.

A unidirectional path is one that contains at least one unidirectional link. The *mutually unidirectionally reachable nodes* metric represents the number of nodes to which a node's shortest paths are unidirectional, whose shortest paths to it are also unidirectional. This metric reflects the likelihood that the routing protocol would encounter a unidirectional path when trying to route between a pair of nodes.

The *one-way unidirectionally reachable nodes* metric represents the number of nodes to which a node's shortest path is unidirectional, but which do not have a route to this node. Only protocols that do not involve two-way per-hop or end-to-end communication can route to destination nodes that are only reachable in one direction (e.g., a protocol using network-wide broadcast to deliver data).

4.3. Path characteristics metrics

The *unidirectional links per path fraction* metric represents the average number of unidirectional links as a fraction of all links in a unidirectional path (over all unidirectional shortest paths in the network). This metric reflects the "unidirectionality" of a unidirectional path, i.e., the level of effort that the routing protocol would need to expend on average when it encounters a unidirectional path. The higher the value for this metric, in general the higher the routing overhead, as each additional unidirectional link may incur additional overhead, for example in terms of multihop acknowledgments.

Path length is an important network routing characteristic, because packet latency, as well as the likelihood of link breaks and packet collisions along a path, are in part proportional to the length of the path. In addition, the discovery and maintenance (monitoring for link breaks) of a longer path may incur higher latency and a higher number of control packet transmissions.

The *average shortest-path length* metric represents the average path length computed over all shortest paths in the network. The *maximum shortest-path length* metric represents the longest path, among all shortest-paths between pairs of nodes, encountered during the lifetime of the network.

The *path length benefit* metric represents the number of hops on average by which the average unidirectional shortest path is shorter than the shortest bidirectional path between a pair of nodes. This metric is computed when the shortest unidirectional path is shorter than any existing bidirectional paths.

4.4. Link and path change metrics

The *link changes* metric represents the average number of link changes per second in the network; a link can change between being reachable and unreachable, and between being unidirectional and bidirectional. The number of link changes reflects the level of mobility in the network, which in turn affects routing protocol performance, as it takes time, and often additional overhead, to react to changes in link directionality and reachability, and to avoid or reduce interruptions in the flow of traffic caused by link changes.

The *unidirectional to bidirectional link changes* metric represents the number of times per second that a unidirectional link becomes bidirectional. Protocols that can detect this condition can automatically reduce their control overhead, as they can stop sending acknowledgment packets across multihop reverse paths but start sending them directly.

The *bidirectional to unidirectional link changes* metric represents the number of times per second that a bidirectional link becomes unidirectional. This metric indicates how many perceived link breaks are actually changes to unidirectional links. Protocols that can distinguish between a link that has become unidirectional and one that has become disconnected can continue routing along the route containing this link, whereas ones that cannot make this distinction would initiate route repair procedures and incur more overhead and potentially a disruption in packet delivery.

The *unidirectional to unreachable link changes* metric represents the number of times per second that a unidirectional link becomes unreachable. This metric reflects the frequency of disconnection in the network, and can also be used in conjunction with the unidirectional to bidirectional link changes metric to indicate whether unidirectional links in a given scenario are more likely to become bidirectional or unreachable.

Mobility and the resulting link changes affect the paths along which packets are forwarded towards their destinations, and may cause the routing protocol to discover alternate routes. The *number of shortest-path changes per second* metric represents the average number of times that a shortest route between two nodes breaks or a shorter path becomes available.

5. Results and analysis

In this section, we analyze the routing characteristics of unidirectional network scenarios produced by the random-, two-, and three-power models with the help of the metrics we defined in Section 4.

5.1. Overview

The mean values for the metrics characterizing the generated network scenarios prior to introducing variations in power (i.e., all nodes have the nominal transmission range) are shown in Table 1. The notable differences in the values of the metrics between scenarios with different levels of mobility are in the total number of neighbors, number of link changes per second, and number of shortest-path changes per second. The total number of neighbors is higher when the network is mobile, due to the fact that the Random Waypoint movement leads to a non-uniform distribution of the nodes, with the density of the network being higher towards the center of the area [13]. The total number of neighbors is similar at different speeds, because the speed of node

Table 1
Network scenarios without power variations

Metric	Static	1 m/s	20 m/s
Total number of neighbors	16.12	22.08	23.24
Unreachable nodes per node	0	0.16	0.03
Average shortest-path length	2.92	2.41	2.34
Maximum shortest-path length	7.1	7	7.5
# Link changes per second	N/A	3.41	43.1
# Shortest-path changes per second	N/A	42.7	496.2

movement in the Random Waypoint model does not affect node distribution and network density at a pause time of 0 [13]. Higher speeds produce a higher number of link and shortest-path changes with the following relationship: a 20-fold increase in speed of motion leads to a factor of 12.6 increase in the number of link changes per second, and a factor of 11.6 increase in the number of shortest-path changes per second. The increase in number of shortest-path changes at a higher speeds is smaller than the increase in the number of link changes, because not all link changes lead to a path change.

In Sections 5.2–5.5, we consider a Random Waypoint movement model with a maximum speed of 1 m/s. Section 5.6 discusses the effects of increasing the speed to 20 m/s, as well as the static network case.

Each point (or bar) in the graphs presenting the results is the average of 10 simulation scenarios for the given pause time and maximum speed, generated as described in Section 3.1. All metrics are averages computed over all nodes and over the entire duration of the simulated scenario (e.g., 900 s).

The axes of the three-dimensional figures presented in this section are not always oriented the same way. This was necessary in order to make visible important features of the graphed data.

In this section, we aim to address the following general questions:

- What is the range of values of each metric for a given model? For example, for a given network density, is there a parameterization of the random-power model that will result in scenarios in which more than 70% of the links in the network are unidirectional? Answering such questions enables protocol designers to select a model and a set of parameterizations of it that would produce scenarios with the desired routing characteristics and level of difficulty.
- What is the relationship between the values of different metrics for a given network density in the context of a given model? For example, is the value of one metric predictive of the values of other metrics?
- Given a unidirectional network scenario, what are its routing characteristics? What are the factors that influence these characteristics? Answering such questions will give researchers insight into the behavior of the network and the effect of this behavior on routing protocol performance. Understanding unidirectional networks would also be useful in designing routing mechanisms and protocols that route over unidirectional links.

5.2. Neighbor-related metrics

The neighbor-related metrics reveal the level of connectivity and average number of unidirectional links in the network, as well as the likelihood that the routing protocol would encounter a unidirectional link (Section 4.1).

5.2.1. Random-power model

As the minimum wireless range, $R_{\min}$, increases, the average wireless range in the network increases as well. As a result, more nodes can reach each other (i.e., more links are formed) and the network becomes better connected, as reflected in the monotonic increase in the total number of neighbors per node (Fig. 2). In addition, for larger values of $R_{\min}$, the differences between the ranges of the nodes become smaller, which leads to a decrease in the likelihood of a link between two nodes being unidirectional. The smaller number of unidirectional links at a larger average range in the network, leads to a decrease in the number of in-neighbors and an increase in the number of bi-neighbors of a node (Fig. 2). The increase in the number of bi-neighbors is faster than the decrease in the number of in-neighbors, because the formation of one bidirectional link creates one link from the point view of each neighbor (i.e., a total

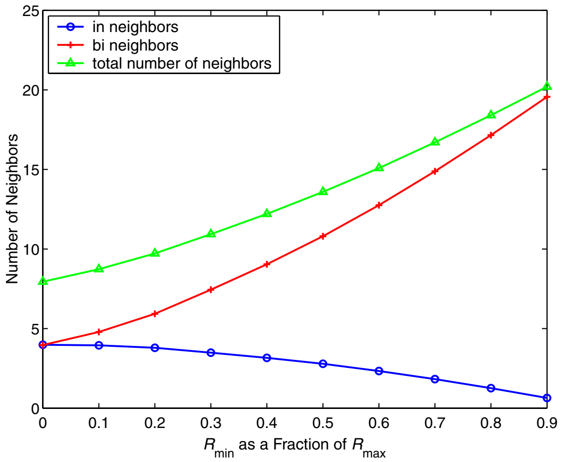

Fig. 2. Random-power model: total number of neighbors per node (1 m/s).

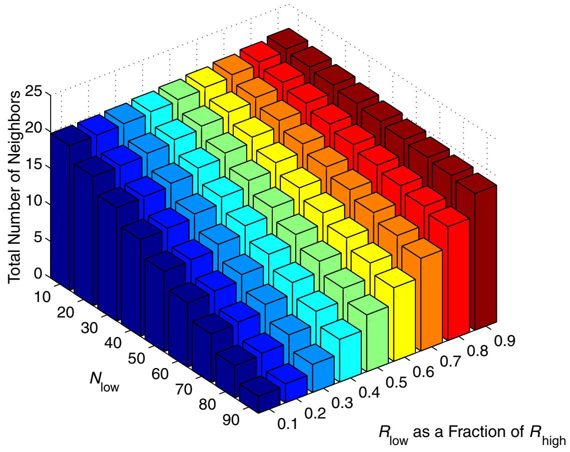

Fig. 3. Two-power model: total number of neighbors per node (1 m/s).

of two one-way links), while the creation or loss of one unidirectional link represents only one one-way link.

In the random-power model, the likelihood of a neighbor being unidirectional reaches a maximum of 50% for R_{min} of 0 m, and decreases monotonically by about 2–5%, down to 3%, as R_{min} increases. An approximation of these results can be obtained analytically: in the random-power model, at R_{min} of 0 m and a maximum range of 250 m, half of the nodes have a range smaller than the mean range (i.e., a range less than 125 m) and half have a range larger than the mean range; thus, on average, half of the links in the network should be unidirectional, as a node is equally likely to have a link with a unidirectional neighbor as it is to have a link with a bidirectional neighbor.

5.2.2. Two-power model

Unlike the random-power model, in which one parameter (the minimum range) changed between scenarios, in the two-power model, two parameters and their interaction determine the routing characteristics of the network; these parameters are the number of low power nodes, or N_{low}, and the wireless range of the low power nodes, R_{low} (Section 3.2). Similarly to the random-power model, the total number of neighbors in the two-power model increases monotonically with the increase in average range (i.e., with the decrease in N_{low} and the increase in R_{low}), and also with the decrease in the differences between the high and low power ranges, R_{high} and R_{low} (Fig. 3). The increase in the total number of neighbors is larger for higher average ranges, because the number of links grows faster when the connectivity in the network is higher and the average range is higher.

The total number of neighbors increases faster with the decrease in N_{low} than with the increase in R_{low}, indicating that the effect of the number of low power nodes is stronger than the effect of the magnitude of their range, even though the average range in the network is equally affected by both parameters. The reason for this is that N_{low} has a stronger effect than R_{low} on the in-neighbors component of the total number of neighbors (Fig. 4): unlike the number of in-neighbors in the random-power model scenarios, which depends on the average range in the network, in the two-power model, the number of in-neighbors is more dependent on the number of nodes with different ranges (i.e., the pairs of nodes which can form unidirectional links) and to a lesser extent on the magnitude of the differences between these ranges. As a result, the number of in-neighbors increases with an increase in N_{low} up to $N_{low} = N_{high}$, which is the point at which the number of pairs of nodes with different ranges is highest, and when N_{low} becomes larger than N_{high}, the

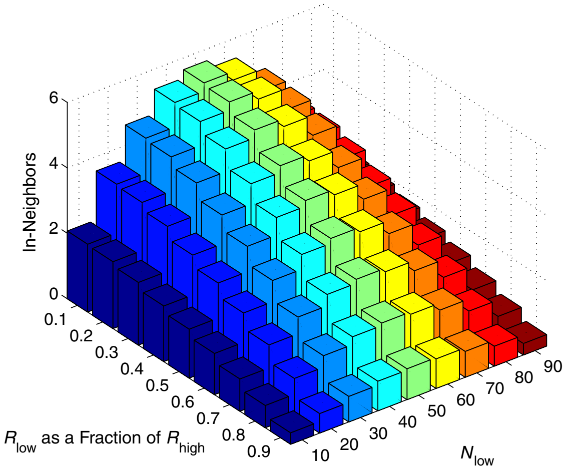

Fig. 4. Two-power model: number of in-neighbors per node (1 m/s).

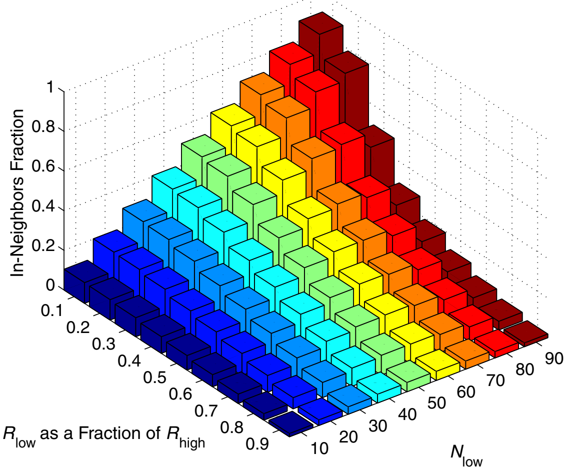

Fig. 5. Two-power model: in-neighbors as a fraction of the total number of neighbors (1 m/s).

number of pairs of nodes with different ranges begins to decrease, which leads to a decrease in the number of in-neighbors as well.

The values for maximum number of total, in-, and bi-neighbors are similar to the corresponding maximum values in the random-power model, whereas the minimum values in the two-power model reach values that are about four times smaller; the two-power model produces scenarios with a wider range of these particular routing characteristics than the random-power model.

Unlike the random-power model, the number of in-neighbors and the fraction of in-neighbors metrics in the two-power model follow different trends: the number of in-neighbors metric achieves its maximum at $N_{low} = 50$ and $R_{low} = 0.1 \times R_{high}$, whereas the in-neighbors fraction metric achieves its maximum at $N_{low} = 90$ and $R_{low} = 0.1 \times R_{max}$ (Fig. 5). The in-neighbors fraction increases monotonically as R_{low} decreases, and for values of R_{low} of 0.1 and 0.2, it increases monotonically with an increase in N_{low} as well. For values of R_{low} higher than 0.2, however, the behavior becomes more complex, as the increase in the in-neighbors fraction stops and turns into a decrease as N_{low} increases. The point at which this switch occurs is different for different R_{low} values, and as R_{low} increases, the switch happens for increasingly lower values of N_{low}. This is due to the fact that for lower values of R_{low}, the effect of the magnitude of the differences in range overwhelms the effect of the number of nodes with different ranges, whereas for higher values of R_{low}, the magnitude of the differences in range decreases, and thus the dominant effect is that of the number of low power nodes, which results in the maximum in-neighbors fraction for each value of R_{low} occurring at an increasingly smaller value of N_{low}.

5.2.3. Three-power model

Similarly to the random- and two-power models, network connectivity in the three-power model, as expressed by the total number of neighbors, increases with the increase in average range, i.e., with the decrease in the number of low and medium power nodes, N_{low} and N_{medium}, respectively. The total number of neighbors is generally higher for Scenario 1, because the average range in this scenario is higher and therefore reachability is higher.

The number of in-neighbors in both Scenarios 1 and 2 increases with the increase in N_{low} and N_{medium} up to a point and then decreases (Scenario 1 is shown in Fig. 6; the Scenario 2 graph, not shown, is similar). The number of in-neighbors metric is more dependent on the value of N_{low} than on the value of N_{medium}, indicated by a higher slope of the graph as N_{low} increases than when

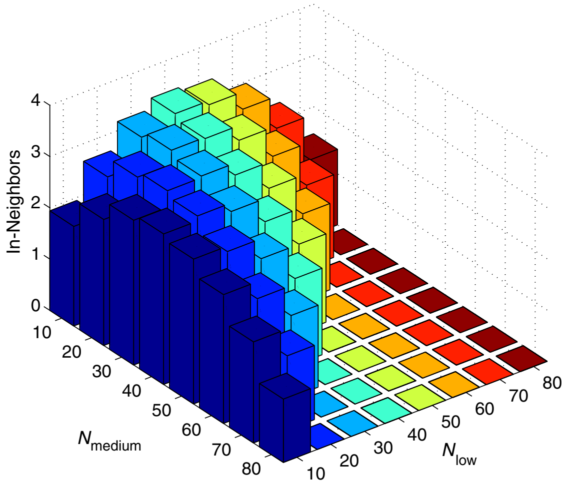

Fig. 6. Three-power model (Scenario 1): number of in-neighbors per node (1 m/s).

N_{medium} increases. This is due to the fact that the low power nodes have larger differences in range with the high power nodes on average, and are thus more likely than the medium power nodes to be involved in unidirectional links with the high power nodes. The values for the number of in-neighbors metric in Scenario 2 are higher by 1.55 on average, whereas the values for the number of bi-neighbors metric are lower by 6.86 on average, as the network is much less connected than in Scenario 2 and the average range is lower.

The maximum value for the number of in-neighbors in both Scenarios 1 and 2 occurs at $(N_{low}, N_{medium}, N_{high}) = (40, 10, 50)$. For each successive value of N_{low}, the highest value of the in-neighbors metric occurs at a smaller value of N_{medium}. In addition, up to the maximum value of the number of in-neighbors for each value of N_{low}, N_{high} is 50. This indicates that the effect of the magnitude of the differences between the ranges of different nodes has a higher impact on the number of unidirectional links in these parameterizations of the three-power model than does the number of such differences (in contrast to the two-power model). Since the highest magnitude of the differences in range occurs when a low power node forms a link with a high power node, the highest value for the number of in-neighbors is achieved when the number of low power nodes and high power nodes are as similar as possible, which happens at point $(40, 10, 50)$. The reason why N_{high} has a higher value than N_{low} at the maximum point is that this is the case in which the highest average magnitude of differences in ranges is created, as the difference between R_{high} and R_{medium} is larger than the difference between R_{medium} and R_{low} (Section 3.2.3).

The values of the in-neighbors fraction metric in Scenario 1 follow the trend of the values of the in-neighbors metric except that the in-neighbors fraction metric reaches its maximum of 0.23 at $(N_{low}, N_{medium}, N_{high}) = (50, 10, 40)$ rather than at $(40, 10, 50)$, the point at which the number of in-neighbors is highest. The maximum is not at $N_{high} = 50$, since at this value, the number of bidirectional links is high as well (these are bidirectional links between the high power nodes) and compensates for the higher number of unidirectional links. As in the two-power model, a higher number of unidirectional links does not necessarily result in a higher fraction of unidirectional links.

Unlike Scenario 1 but similarly to the two-power model, the in-neighbors fraction metric in Scenario 2 exhibits a different behavior from the in-neighbors metric (Fig. 7). It achieves values of up to 0.8, which is about 3.4 times higher than in Scenario 1, and increases monotonically with successively higher values of N_{low}, reaching its maximum value at (80, 10, 10). Since the ranges in

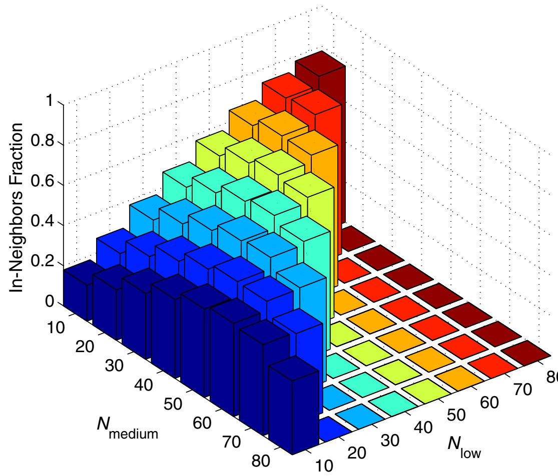

Fig. 7. Three-power model (Scenario 2): in-neighbors as a fraction of the total number of neighbors (1 m/s).

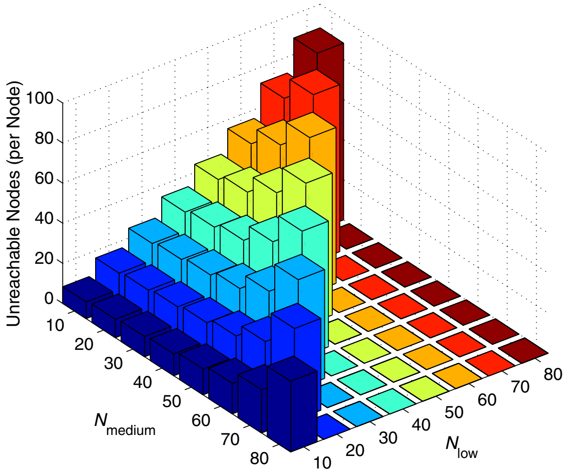

Fig. 8. Three-power model (Scenario 2): number of nodes unreachable from a node (1 m/s).

Scenario 2 are lower and the differences between them higher than in Scenario 1 (Section 3.2.3), the level of unreachability is high (Fig. 8), and as a result the in-neighbors fraction is also high: when the number of low power nodes is high in Scenario 2, pairs of nodes that are reachable usually include a low power node and a medium or high power node (pairs of low power nodes are likely to be unreachable) and such pairs are likely to be connected via a unidirectional link due to the differences in range.

5.2.4. Summary

The two- and three-power models can produce unidirectional network scenarios with a larger range of values for the neighbor-related metrics than the random-power model. In general, models that have a larger number of configurable parameters can achieve a larger set of routing characteristics, but are also more complex to analyze. For example, in the random-power model, the number of in-neighbors and the fraction of in-neighbors follow the same trend, as average range increases across the set of unidirectional scenarios (Section 5.2.1), whereas in the two- and three-power models, the number of in-neighbors and the in-neighbors fraction may exhibit different behaviors from each other (Section 5.2.2). In addition, some metrics do not exhibit a uniform behavior in scenarios produced by successive parameter values of a given model; for example, the in-neighbors fraction in the two-power model exhibits a non-uniform behavior for different R_{low} values, achieving its maximum values for different values of N_{low}. These findings indicate that choosing a range of values for parameterizing a model for generating unidirectional links and expecting that the likelihood that a routing protocol would encounter unidirectional links increases monotonically with each successive parameterization is not necessarily a good strategy. The routing characteristics of the scenarios produced by each parameterization need to be analyzed in detail.

5.3. Node reachability metrics

The node reachability metrics (Section 4.2) reflect the likelihood that a node can reach an arbitrary node in the network (i.e., the likelihood of partitions), and the type of reachability it would have with that arbitrary node (i.e., via a bidirectional, mutually unidirectional, or one-way unidirectional path). The higher the likelihood of encountering a unidirectional path, the higher the overhead the routing protocol would have to expend due to having to route over unidirectional links, or due to trying to find alternate bidirectional paths (Section 4.2).

5.3.1. Random-power model

The number of unreachable nodes per node starts out high (30% of all nodes) at $R_{\min}$ of 0 m (Fig. 9). As the average range in the network increases, previously unreachable nodes become unidirectional or bidirectional neighbors, and as a result, the number of unreachable nodes drops sharply and becomes negligible at $R_{\min}$ of 0.3–0.4. A significant number of nodes are reachable only one-way at $R_{\min}$ of 0 m as well. The similarity between the values of the unreachable nodes and one-way unidirectionally reachable nodes metrics is due to the fact that poorly connected nodes (e.g., ones close to the edges of the site (Section 5)), flap between being barely connected and being partitioned as they move about. In addition, the smaller the average range, the more likely it is for a node to be partitioned (rather than unidirec-

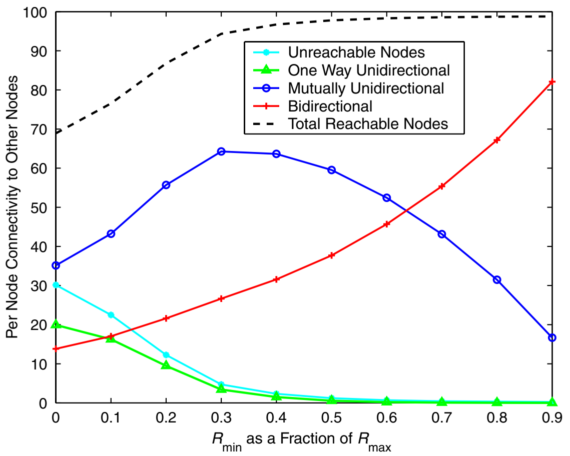

Fig. 9. Random-power model: per node connectivity to other nodes (1 m/s).

tionally connected to the network). As a result the unreachable nodes graph has higher values than the one-way unidirectionally reachable nodes graph, and the differences are more pronounced for lower values of R_{min}.

The number of bidirectionally reachable nodes per node rises monotonically as the average range in the network increases (similarly to the number of bi-neighbors metric, Fig. 2). The number of unidirectionally reachable nodes per node, on the other hand, does not follow the monotonic decrease in the number of in-neighbors but instead rises up to R_{min} of 0.3 and only then begins to decrease. This initial increase is due to the initial increase in connectivity and average wireless range which conceptually draws nodes closer together as a result of which unidirectional links become bidirectional and unreachable links become unidirectional. The number of unidirectionally reachable nodes per node starts to fall when R_{min} becomes larger than 0.3 of R_{max}, since node unreachability becomes negligible and the dominant effect of an increase in average range is that of unidirectional links becoming bidirectional.

The number of unidirectionally reachable nodes is higher than what the average in-neighbors fraction would suggest (Section 5.2.1) because the addition of unidirectional links has a stronger impact on the unidirectional paths metric than on the in-neighbors fraction metric; the addition of one unidirectional link can make multiple paths unidirectional and vice versa, the removal of one unidirectional link may make multiple paths bidirectional.

5.3.2. Two-power model

Unreachability increases with a decrease in average range, and at $N_{low} = 90$ and $R_{low} = 0.1 \times R_{high}$, it reaches a value of 98% (Fig. 10), which is over three times the maximum number of unreachable nodes per node in the random-power model. Unlike the random-power model, the one-way unidirectionally reachable nodes metric (Fig. 11) does not follow the same trend as the unreachable nodes metric because in addition to average range, it is also influenced by the number of pairs of nodes with differing ranges; in the random-power model both of these parameters change together as only one parameter (the minimum range) is varied. As a result, the maximum value for the one-way unidirectionally reachable nodes metric occurs at the same values of R_{low} and N_{low} as the highest value of the in-neighbors metric ($R_{low} = 0.1$ and $N_{low} = 50$).

The number of bidirectionally reachable nodes per node increases monotonically with N_{low}, for values of R_{low} of 0.1 and 0.2 (Fig. 12). For larger values of R_{low} it decreases with an increase in N_{low}

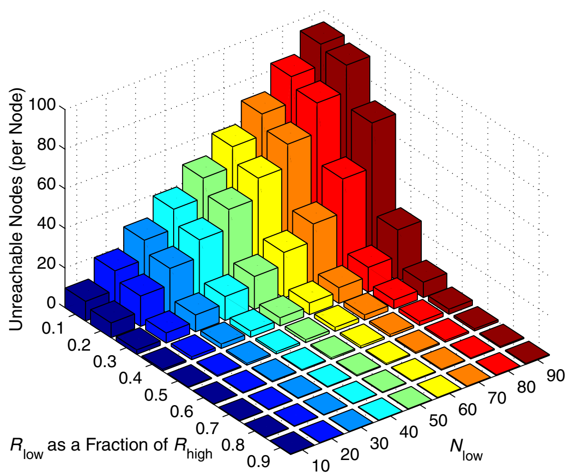

Fig. 10. Two-power model: number of nodes unreachable from a node (1 m/s).

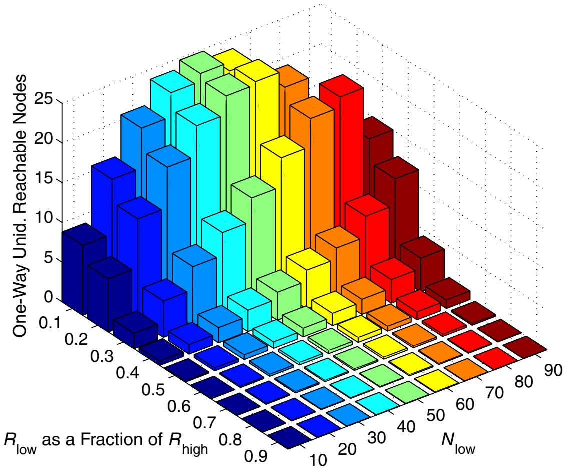

Fig. 11. Two-power model: number of one-way unidirectionally reachable nodes per node (1 m/s).

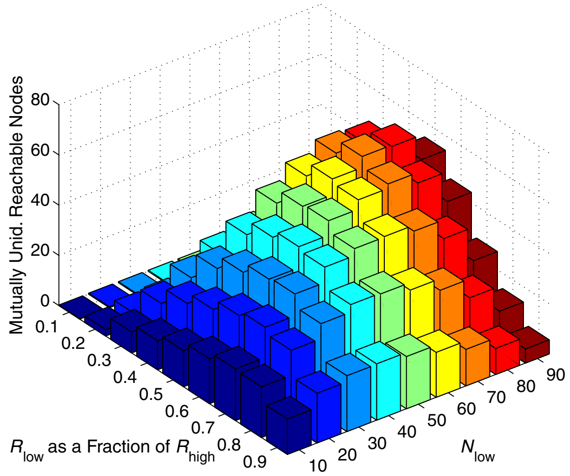

Fig. 13. Two-power model: number of mutually unidirectionally reachable nodes per node (1 m/s).

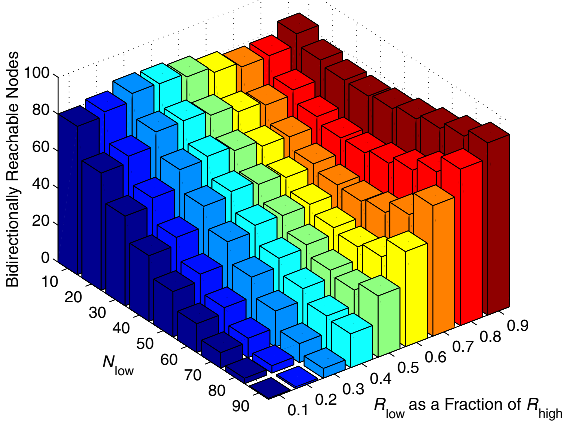

Fig. 12. Two-power model: number of bidirectionally reachable nodes per node (1 m/s).

up to a point, and then starts to increase since N_{low} becomes higher than N_{high} after which point the majority of bidirectional paths are between low power nodes (since R_{low} is high) and these paths dominate the overall shortest-path reachability. As R_{low} increases, the number of bidirectional paths increases as well, except that for values of N_{low} smaller than 40, there is a temporary decrease in the number of bidirectional paths around R_{low} of 0.6 and 0.7. This decrease is due to the fact that at some point, nodes that were previously unreachable are now reachable due to the higher average range, but they are reachable only unidirectionally, as the higher range is not high enough to make them bidirectional neighbors.

The number of mutually unidirectionally reachable nodes metric behaves differently from the number of in-neighbors metric. It increases with increasing values of N_{low} and R_{low} (for low N_{low} and R_{low} values), and then begins to decrease (Fig. 13). The initial increase is due to increased connectivity in a poorly connected network in the case of N_{low}, and an increased number of pairs of nodes with differing ranges in the case of R_{low}. The subsequent decrease is due to a decrease in the number of unidirectional links.

5.3.3. Three-power model

There is very little unreachability in Scenario 1 (a maximum of 2.7 unreachable nodes per node; figure not shown). The number of unreachable nodes per node rises with a decrease in average range, i.e., as N_{low} and N_{medium} increase. Similarly to the random-power model, the number of one-way unidirectionally reachable nodes graph tracks the unreachable nodes graph very closely and has slightly lower values.

In Scenario 2, unreachability also increases with both an increase in N_{low} and N_{medium}, except, due to the smaller values of R_{low} and R_{medium}, the

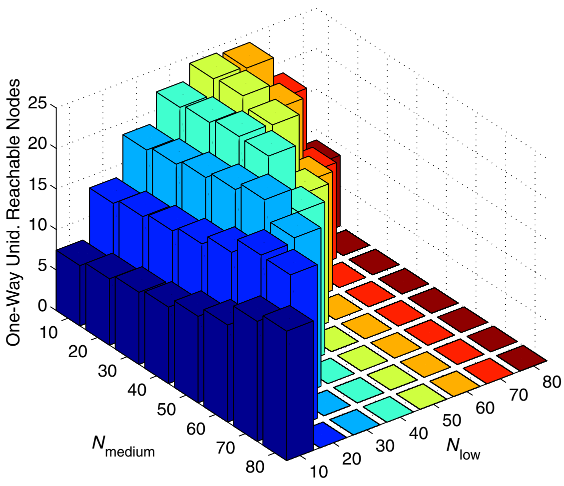

Fig. 14. Three-power model (Scenario 2): number of one-way unidirectionally reachable nodes per node (1 m/s).

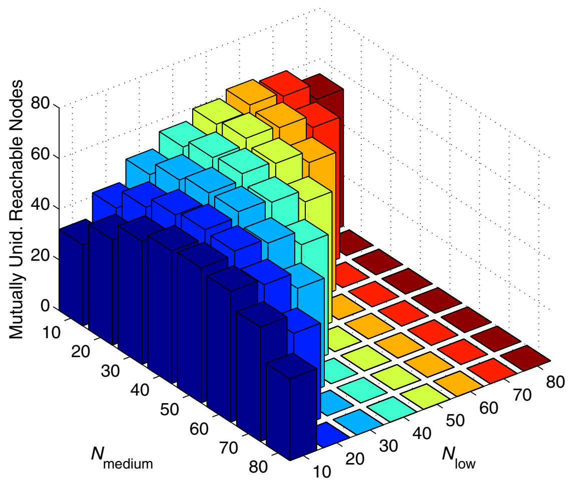

Fig. 15. Three-power model (Scenario 1): number of mutually unidirectionally reachable nodes per node (1 m/s).

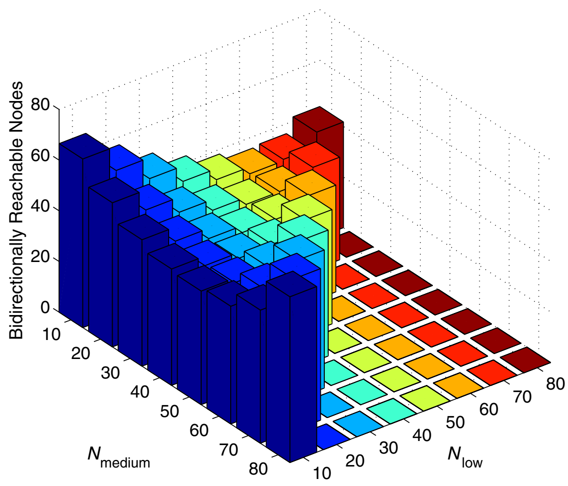

Fig. 16. Three-power model (Scenario 1): number of bidirectionally reachable nodes per node (1 m/s).

unreachability is much higher than in Scenario 1, reaching up to 90 unreachable nodes per node (i.e., 91% of the nodes) (Fig. 8). Unlike Scenario 1, in Scenario 2 the number of one-way unidirectionally reachable nodes initially increases with the increase in average range, and then starts to decrease along with the number of unreachable nodes due to the increasing connectivity in the network which leads to the formation of a higher number of mutually unidirectionally reachable and bidirectionally reachable nodes (Figs. 14 and 8).

The mutually unidirectionally reachable nodes and the bidirectionally reachable nodes' graphs in Scenario 1 follow complementary trends (Figs. 15 and 16) as virtually all paths are either bidirectional or mutually unidirectional due to the high level of connectivity. In Scenario 2, the bidirectionally reachable nodes graph follows a complementary trend to the unreachable nodes graph (Fig. 8) as nodes are more likely to be unreachable or bidirectionally connected to the network than to be unidirectionally connected to the network, due to the large differences between the ranges of different nodes. The number of mutually unidirectionally reachable nodes in Scenario 2, increases with an increase in N_{medium} as the medium power nodes have a stronger impact on the growth of the number of unidirectional paths in the network than do the low power nodes, since the medium power nodes are more likely to participate in unidirectional links than the low power nodes which are more likely to be unreachable (due to their small ranges). Unlike in the other power models and Scenario 1, the connectivity of the network remains low even in the most connected configuration of Scenario 2 and as a result, the number of mutually unidirectionally reachable paths does not reach a peak value and then top off but only increases in our experiments. The in-neighbors fraction actually does top off (Fig. 7) but the number of unidirectional links in the network has a weaker effect on the composition of paths than on the fraction of unidirectional links, as one unidirectional link can cause multiple paths to be unidirectional.

5.3.4. Summary

The analysis of the node reachability metrics shows that attempting to predict the routing characteristics of the network based on knowledge of only the number or fraction of unidirectional links in the network can be misleading as the values of these metrics are not indicative of the types of paths that the routing protocol is likely to encounter. In addition, picking seemingly similar parameterizations can lead to network scenarios with very different characteristics, and vice versa, picking seemingly different parameterizations can lead to network scenarios with similar routing characteristics. Understanding the reachability characteristics of the network would enable protocol designers to better analyze routing protocol behavior, as different routing mechanisms are sensitive to the presence of different kinds of paths. Similarly to the neighbor-related metrics (Section 5.2), the two- and three-power models provide a wider range of values for the reachability metrics but are more complex to analyze than the random-power model.

5.4. Path characteristics metrics

The path characteristics metrics (Section 4.3) characterize a scenario in terms of the average and maximum shortest-path length in the network, the path length benefit of using a shorter unidirectional path instead of a longer bidirectional path, and the level of unidirectionality of a unidirectional path, i.e., the fraction of links on each unidirectional path that are unidirectional. Only shortest paths are discussed in this section as these are the paths the routing protocol is most likely to use.

5.4.1. Random-power model

As the average range in the network increases, so does the number of bidirectional paths. As a result, the likelihood of encountering a bidirectional path that is as short as the corresponding unidirectional path between a pair of nodes increases also. This trend is reflected in the path length benefit of using unidirectional links which decreases from 1.4 to 1 and remains equal to 1 for values of $R_{\min}$ greater than 0.6, which is the point at which the dominant reachability between nodes in the network starts to be via bidirectional paths (Fig. 9). The path length benefit is largely insignificant in the random-power model.

The average and maximum shortest-path lengths in the network decrease as connectivity increases and more paths become available. The maximum shortest-path length varies between 14.8 and 7.3 for $R_{\min} = 0$ and $R_{\min} = 0.9 \times R_{\max}$, respectively, whereas the average shortest-path length varies between 3.6 and 2.3 for the same $R_{\min}$ values. The decrease in average shortest-path length is small due to the dominant contribution of 1- and 2-hop paths, whose number increases the fastest with increased connectivity.

The unidirectional links per path fraction is about 45% at $R_{\min}$ of 0 m and decreases monotonically to about 38%. There are several competing factors that influence this metric. First, as $R_{\min}$ increases, there are less unidirectional links and since links are shared between paths, the conversion of a unidirectional link to a bidirectional link (as a result of the higher average range) may affect more than one path (leading to a lower unidirectional links per path fraction). Another factor that influences this metric is that path length decreases with an increase in $R_{\min}$ and shorter paths do not share as many unidirectional links because there are less nodes at which they can intersect (leading to a lower unidirectional links per path fraction). Finally, shorter paths can have a higher fraction of unidirectional links with a smaller number of unidirectional links on them (which leads to a higher unidirectional links per path fraction). The combination of these factors leads to a slow decrease in the unidirectional links per path fraction.

5.4.2. Two-power model

The average path length benefit of using unidirectional paths reaches a maximum of 2.16 hops, which is higher than in the random-power model. This is due to the wider range of values for the number of unreachable and unidirectionally reachable nodes in the two-power scenarios (Section 5.3.2). As a result, at low values of N_{low} and R_{low}, where the number of unidirectional links is high, longer bidirectional paths are less likely to exist than shorter ones; a single unidirectional link can make the path unidirectional, and the more links in a path, the more likely it is that one of them may be unidirectional. As in the random-power model, the path-length benefit declines with the increase in the number of bidirectional paths.

The wider range of values for the unreachable and unidirectionally reachable nodes (relative to the random-power model), leads to a wider range for the shortest and maximum shortest-path metrics as well, which vary between 1.87 and 4.57, and 7 and 21.3, respectively. The range of values for the unidirectional links per path metric is also wider in the two-power model (0.12–0.59). Similarly to the random-power model, this metric is highest for scenarios with the highest in-neighbors fraction and lowest for scenarios with the lowest in-neighbors fraction, $(N_{\mathrm{low}}, R_{\mathrm{low}}) = (90, 0.1)$ and $(10, 0.9)$, respectively.

5.4.3. Three-power model

The path length benefit in both Scenarios 1 and 2 decreases with the increase in the number of bidirectional paths, just as in the random- and two-power models. In Scenario 1, the maximum path length benefit is only 1.22, because the number of unreachable nodes is very low (Section 5.3.3). The fraction of nodes that a node cannot reach on average reaches 91% in Scenario 2 (Section 5.3.3) leading to a higher path length benefit of using unidirectional links (up to 2.9 hops).

Similarly to the random- and two-power models, path length increases with an increase in unreachability, for both Scenarios 1 and 2. The average shortest-path length and maximum shortest-path length reach values of 3.57 and 12.7 in Scenario 1, and 4.45 and 20.8 in Scenario 2, respectively.

The unidirectional links per path fraction metric achieves ranges of 0.34–0.42 for Scenario 1 and 0.33–0.57 for Scenario 2, which are wider than the ones in the random-power model but narrower than the ones in the two-power model.

5.4.4. Summary

The analysis of the path characteristics metrics reveals that unidirectional links can have a significant impact on the length of network routes, and that unidirectional paths in the three-power models generally contain a significant fraction of unidirectional links. The path length benefit of using unidirectional paths is fairly small though of course, in some cases bidirectional paths are actually not available, so protocols that attempt to find a bidirectional path and are not able to route over unidirectional paths, are not going to be able to deliver their data and may incur unnecessary overhead.

As in the case of the neighbor-related and reachability metrics, the two- and three-power models provide a larger set of values for each metric than the random-power model and thus provide protocol designers with more choices for experimenting with routing protocols in unidirectional networks with different characteristics.

5.5. Link and path changes metrics

The link and shortest-path changes metrics reflect the level of mobility in the network and the challenge to the routing protocol in terms of distinguishing between changes in the directionality of a link versus a link break, which affects protocol efficiency (Section 4).

5.5.1. Random-power model

The number of links in the network increases with an increase in average range and as a result, the number of link changes increases as well (Fig. 17); node movement in the presence of a higher number of links causes a higher number of link changes. However, the fraction of all links that experience a change actually decreases from 0.009 at R_{min} of 0 to 0.0065 at R_{min} of 0.9 of R_{max}. This decrease is due to the fact that the rise in the number of links outpaces the increase in the number of link changes, since at a higher average range, the distance traversed and thus the time required to cause a link change is longer.

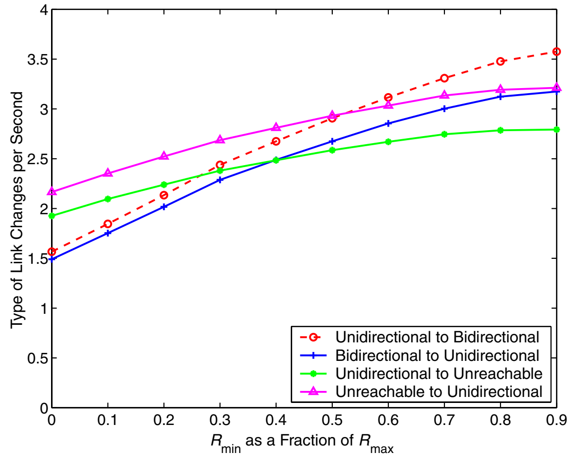

Fig. 17. Random-power model: link changes (1 m/s).

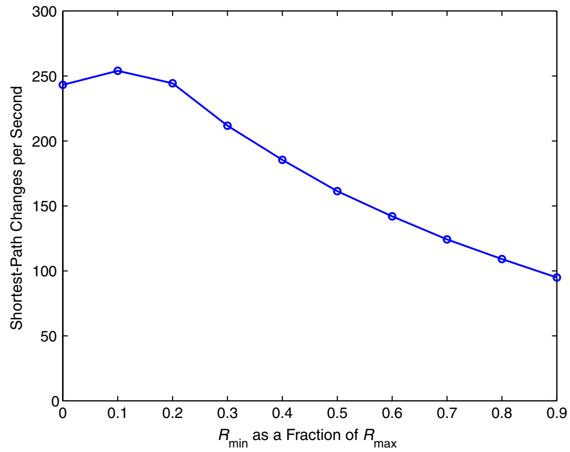

Fig. 18. Random-power model: shortest-path changes (1 m/s).

The number of shortest-path changes and the fraction of shortest paths that change decrease with an increase in average range (Fig. 18), with the fraction of shortest-path changes decreasing from 0.035 to 0.01 at $R_{\min} = 0$ and $R_{\min} = 0.9 \times R_{\max}$, respectively. This is due to the fact that as average range increases, connectivity increases leading to the presence of redundant links (i.e., links that do not improve reachability or path length) and when a node encounters such a link or when such a link breaks, the shortest paths between it and other nodes are unaffected. This effect is reflected in the ratio of the fraction of shortest-path changes to the fraction of link changes which starts out at 3.9 and decreases monotonically, reaching a value of 1.51 at $R_{\min}$ of 0.9 of $R_{\max}$.

We have divided link changes into several groups to highlight several different types of link-related events that may affect the routing protocol (Section 4). The only statistically significant differences between the four types of link changes are for the smallest and largest values of $R_{\min}$ (Fig. 17). For small values of $R_{\min}$, the likelihood of unidirectional links becoming unreachable and unreachable links becoming unidirectional is higher than the likelihood of unidirectional links becoming bidirectional and bidirectional links becoming unidirectional. For the highest several values of $R_{\min}$, this trend is reversed. These effects are due to the fact that at low values of average range, movement is more likely to cause two nodes to go from being unidirectional neighbors to being disconnected and vice versa, whereas for high values of the average range in the network, movement is more likely to cause them to become bidirectional neighbors.

5.5.2. Two-power model

The number of link changes per second is influenced by the number of links in the network and also by the number of unidirectional links in the network (Fig. 19). The highest number of link changes occurs at the point of highest value for

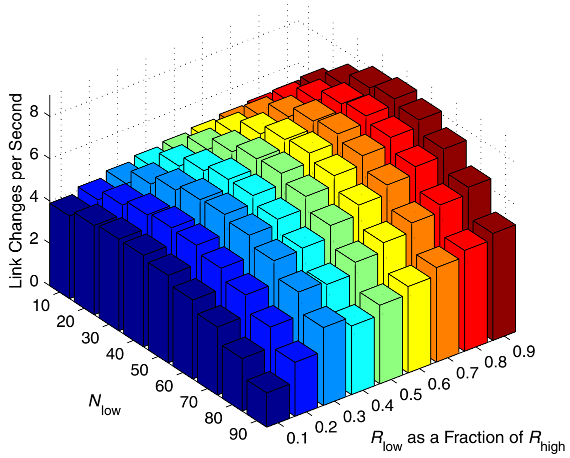

Fig. 19. Two-power model: link changes (1 m/s).

the number of in-neighbors ($N_{low} = 50$) and highest number of links ($R_{low} = 0.9$). The maximum number of link changes per second in the two-power model is 8.16, and the maximum number of link changes per second in the random-power model is 14.17, but the maximum fraction of link changes approaches 0.01 in both models.

The fraction of shortest-path changes reaches a maximum of 0.12 which is 3.43 times higher than the maximum value of the fraction of shortest-path changes in the random-power model. The fact that a similar fraction of link breaks in the two models causes a different fraction of route breaks is due to the differences in path length between the models (Sections 5.4.1 and 5.4.2); longer paths are more affected by link changes than shorter ones.

Similarly to the random-power model, the number of changes between unreachable and unidirectional links is higher than the number of changes between unidirectional and bidirectional links when average range is low and the reverse is true for high average range.

5.5.3. Three-power model

The number of link changes in Scenario 1 increases with average range, e.g., as both N_{low} and N_{medium} increase (Fig. 20), and also with the number of unidirectional links in the network, which is why the link changes curve follows the same trend exhibited by the number of in-neighbors curve. When the number of unidirectional links is highest, the number of link changes is also highest, as unidirectional links participate in the most types of link changes—unidirectional to bidirectional or unidirectional to unreachable and vice versa. Scenario 2 follows the same trend as well, except there are less link changes (about 2, at maximum) since the connectivity is lower as there are less links overall.

Similarly to the random- and two-power models, the number of shortest-path changes is influenced by the lengths of the paths in the network. The shortest-path changes in Scenario 1 increase monotonically with both N_{low} and N_{medium} following the average shortest-path length trend. The number of shortest-path changes in Scenario 2 (Fig. 21) follows the same pattern except that it has higher values than Scenario 1 up to N_{low} of 50 (due to the much lower connectivity) and after that starts to decline (rather than increase monotonically as in Scenario 1) due to the higher level of unreachability (Fig. 8). The fraction of link changes reaches the same maximum value of 0.01 in both Scenarios as in the random- and two-power models, whereas the shortest-path changes fraction in Scenario 2 reaches a maximum of 0.08 (which is smaller than in the two-power model

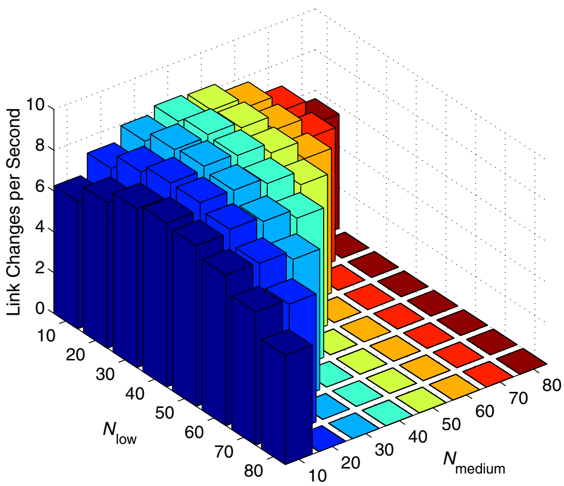

Fig. 20. Three-power model (Scenario 1): link changes (1 m/s).

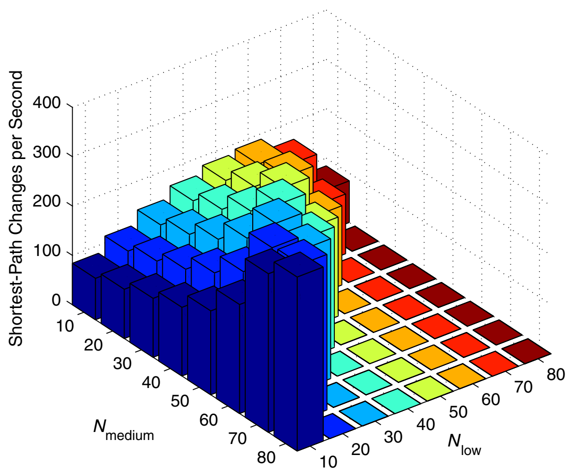

Fig. 21. Three-power model (Scenario 2): shortest-path changes (1 m/s).

and larger than in the random-power model) and is about 4 times higher than the maximum value in Scenario 1 (which itself is higher than the maximum value of the metric in the random-power model). As mentioned in Section 5.5.2, a similar fraction of link breaks in all of the models causes a different fraction of shortest-path changes in each model due to the differences in path length between scenarios generated by each model (Sections 5.4.1, 5.4.2, and 5.4.3); longer paths are more affected by link changes than shorter ones.

The different types of link changes in Scenario 1 have similar values to each other and follow the same trend as the link changes metric.

In Scenario 2, due to the higher level of unreachability, the unidirectional to unreachable link changes have higher values than the unidirectional to bidirectional changes and vice versa.

5.5.4. Summary

The analysis of the link and path change metrics shows that knowing the number or fraction of link changes, does not always help in predicting the number and fraction of shortest-path changes, as these are also dependent on path length and the types of paths that exist in the network. Our findings reinforce once again, that it is important to analyze the routing characteristics of the network using a rich set of metrics, as different metrics expose different aspects of the routing environment that often cannot be predicted based on the values of other metrics.

5.6. Effects of mobility and speed of movement

To explore the effects of mobility, we performed the same set of experiments described in Section 5 but with a maximum speed of movement of 20 m/s, instead of 1 m/s. In addition, we repeated the experiments for a static network.

5.6.1. High mobility

The values of most of the metrics for the random-power model at 20 m/s are nearly identical to the values of the metrics in the 1 m/s scenarios because they are influenced by node density rather than mobility, and the Random Waypoint model maintains a similar density between scenarios with different maximum speeds at a pause time of 0 (Section 5). The metrics that are influenced by the level of mobility in the network are the link and shortest-path changes metrics for which the shapes of the graphs are the same but the absolute values are higher at 20 m/s due to the higher speed of movement. The number of link changes is about 12.7 times higher, and the number of shortest-path changes is about 11.5 times higher at 20 m/s, which match the differences in the ratios of link and shortest-path change metrics between 1 and 20 m/s in the scenarios without power variations (Table 1).

The metrics for the two- and three-power models at 20 m/s exhibit the same behavior as in the random-power model.

5.6.2. Static networks

In a static network with random-power model generated power levels, the total number of neighbors is 74% of the total number of neighbors in a mobile scenario, and the total number of paths is 95% that of the mobile scenarios. This lower connectivity is present in scenarios without power variation as well (Table 1) and is due to the difference in the distribution of nodes on the site between a static network and a network in which the nodes move according to a Random Waypoint model (Section 5). The difference in the distribution of nodes on the site, and the consequent difference in connectivity between the static and mobile scenarios is the cause of differences in their routing characteristics. The general trend in the values of the metrics in the static network is the same as in the mobile ones, except that the number of neighbors is lower, the number of unreachable nodes is higher, and the path lengths are higher. These differences are generally in the range of 5–15%.

5.7. Routing example analysis

In this section, we briefly revisit the AODV example introduced in Section 1, which we are now able to analyze using the metrics studied in the preceding sections.

As mentioned in Section 1, the number and fraction of in-neighbors decrease monotonically in the parameterizations of the two-power model,

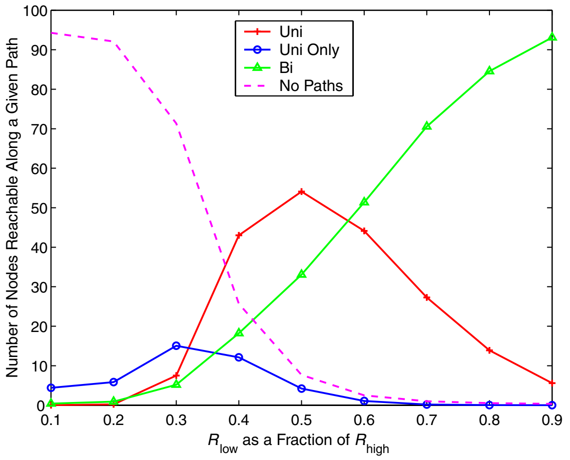

Fig. 22. Two-power model: node reachability metrics for $N_{low} = 90$ (1 m/s).

moving from left to right in Fig. 1. These simple metrics, however, do not explain the non-monotonic behavior of the routing protocol, which can now be interpreted with the help of the metrics we developed. In particular, Fig. 22 shows the reachability metrics for the unidirectional scenarios used in the simulation of AODV.

AODV generates Route Requests when it encounters a unidirectional or broken link, and when a destination node is unreachable. The peak of the Route Request graph is at R_{low} of 0.3; this is the scenario with the highest number of unreachable and one-way unidirectionally reachable paths, and these cause the largest number of route discoveries. For values of R_{low} of 0.1 and 0.2, connectivity is lower than at 0.3, but it is so low that Route Request packets are not able to propagate to many nodes. At R_{low} of 0.4, the number of unreachable nodes is much lower than at 0.3, and even though the number of mutually unidirectional paths is higher (triggering protocol reaction), the number of bidirectional paths is also higher, causing the Route Request curve to begin to decline. This decline turns into a slow increase at R_{low} of 0.6, which is the point at which the number of bidirectional paths in the network begins to exceed the number of unidirectional paths. At first glance, it seems that the number of Route Requests should begin to decline as the protocol is more likely to find bidirectional paths. However, as the number of bidirectional paths grows, so does the number of link breaks as well as the number of bidirectional to unidirectional link changes (Fig. 19). As a result, the protocol is forced to perform more local repairs, both in response to link breaks and also in response to links changing from bidirectional to unidirectional, which are perceived by the protocol as link breaks.

6. Conclusion

In this paper, we have studied the impact of unidirectional links on the routing characteristics of multihop wireless ad hoc networks. Our analysis focuses on *mobile* networks composed of heterogeneous devices with different transmission capabilities. To generate such networks, we used the three most commonly used power variation models for simulations of ad hoc networks with unidirectional links: the random-, two-, and three-power models, each parameterized with a wide range of parameter values. We presented a set of metrics that expose the routing characteristics of the resultant network scenarios, indicating the routing difficulty that each poses to the routing protocol. Our analysis shows that it is important to examine the behavior of the network from multiple viewpoints, as the difficulty and effects of a unidirectional scenario on routing protocol performance can be interpreted only with knowledge of the routing characteristics of the network and insight into their interactions. Our findings enable protocol designers to better choose a set of network scenarios that truly explore a wide range of a routing protocol's behaviors in the presence of unidirectional links, and to better understand the subtleties of the interplay between routing mechanisms and network states.

Acknowledgement

This work was supported in part by NSF under grants CCR-0209204, CNS-0325971, and CNS-0338856, by NASA under grant NAG3-2534, and by a gift from Schlumberger. The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either express or implied, of NSF, NASA, Schlumberger, Carnegie Mellon University, Rice University, or the US Government or any of its agencies.

References

[1] L. Bao, J.J. Garcia-Luna-Aceves, Link-state routing in networks with unidirectional links, in: Proceedings of the IEEE International Conference on Computer Communications and Networks (ICCCN), October 1999, pp. 358–363.

[2] J. Broch, D.A. Maltz, D.B. Johnson, Y.-C. Hu, J.G. Jetcheva, A performance comparison of multi-hop wireless ad hoc network routing protocols, in: Proceedings of the Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking, October 1998, pp. 85–97.

[3] K. Fall, K. Varadhan (Eds.), The *ns* manual (formerly *ns* notes and documentation). The VINT Project, December 2003. Available from: <http://www.isi.edu/nsnam/ns/>.

[4] Z.J. Haas, A Routing Protocol for the reconfigurable wireless network, in: 1997 IEEE 6th International Conference on Universal Personal Communications Record. Bridging the Way to the 21st Century, (ICUPC '97), vol. 2, October 1997, pp. 562–566.

[5] D.B. Johnson, D.A. Maltz, Dynamic source routing in ad hoc wireless networks, in: T. Imielinski, H. Korth (Eds.), Mobile Computing, Kluwer Academic Publishers, Boston, 1996, pp. 153–181 (Chapter 5).

[6] D.B. Johnson, D.A. Maltz, Y.-C. Hu, The Dynamic Source Routing Protocol for Mobile Ad Hoc Networks, Internet-Draft, draft-ietf-manet-dsr-10.txt, July 2004, Work in progress.

[7] M.K. Marina, S.R. Das, Routing performance in the presence of unidirectional links in multihop wireless networks, in: Proceedings of the 2002 ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2002), June 2002, pp. 12–23.

[8] M.R. Pearlman, Z.J. Haas, B.P. Manvell, Using multi-hop acknowledgements to discover and reliably communicate over unidirectional links in ad hoc networks, in: Proceedings of the IEEE Wireless Communications and Networking Conference 2000 (WCNC), September 2000, pp. 532–537.

[9] C.E. Perkins, P. Bhagwat, Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers, in: Proceedings of the SIGCOMM '94 Conference on Communications Architectures, Protocols and Applications, August 1994, pp. 234–244.

[10] C.E. Perkins, E.M. Royer, S.R. Das, Ad Hoc On-Demand Distance Vector (AODV) Routing, RFC 3561, July 2003.

[11] V. Ramasubramanian, R. Chandra, D. Mosse, Providing a bidirectional abstraction for unidirectional ad hoc networks, in: Proceedings of IEEE Infocom, June 2002.

[12] V. Ramasubramanian, D. Mosse, Statistical analysis of connectivity in unidirectional ad hoc networks, in: Proceedings of the International Conference on Parallel Processing Workshops (ICPPW), August 2002, pp. 109–116.

[13] G. Resta, P. Santi, An analysis of the node spatial distribution of the random waypoint mobility model for ad hoc networks, in: Principles of Mobile Computing (POMC) Workshop, October 2002.

[14] P. Sinha, S. Krishnamurthy, S. Dao, Scalable unidirectional routing with Zone Routing Protocol (ZRP) extensions for mobile ad-hoc networks, in: Proceedings of the IEEE Wireless Communications and Networking Conference 2000 (WCNC), September 2000, pp. 1329–1339.

Jorjeta G. Jetcheva is currently a Systems Architect at Firetide, Inc., where she is responsible for driving the evolution of the core mesh routing technology of the company. She has been involved in ad hoc networking research for the past seven years, including the design, analysis, simulation and measurement evaluation of routing protocols and mechanisms. She received her Ph.D. in Computer Science from Carnegie Mellon University's School of Computer Science in May 2004, defending a dissertation on "Adaptive Demand-Driven Multicast Routing in Multi-Hop Mobile Ad Hoc Networks". She has also been involved in Internet traffic measurement and engineering during an internship at Sprint Advanced Technology Laboratories in Burlingame, CA in 2000. Prior to obtaining her Master's and Ph.D. degrees from Carnegie Mellon, she received a B.A. degree in Computer Science and Mathematics, with a summa cum laude in Computer Science, from Mount Holyoke College. While there she was also involved in research on networked workstations, computer vision and computational geometry, and did an internship at the DataViews Corporation, Northampton, MA, where she worked on dynamic data visualization. She has served as a reviewer for numerous networking conferences and journals. She is a member of ACM SIGMOBILE, Sigma Xi and Phi Beta Kappa.

David B. Johnson is an Associate Professor of Computer Science and Electrical and Computer Engineering at Rice University. Prior to joining the faculty at Rice in 2000, he was an Associate Professor of Computer Science at Carnegie Mellon University, where he had been on the faculty for eight years. He is leading the Monarch Project, developing adaptive networking protocols and architectures to allow truly seamless wireless and mobile networking. He has also been

very active in the Internet Engineering Task Force (IETF), the principal protocol standards development body for the Internet, in the Mobile IP and MANET (Mobile Ad Hoc Networks) working groups. He was one of the main designers of the IETF Mobile IP protocol for IPv4 and is the primary designer of Mobile IP for IPv6, and his group's Dynamic Source Routing protocol (DSR) for ad hoc networks has been approved to be published as an Experimental protocol for the Internet. He has served as the General Chair for MobiCom 2003 and the Technical Program Chair for MobiHoc 2002 and MobiCom 1997, has been a member of the Technical Program Committee for over 30 international conferences and workshops, and has been an editor for several journals. He is an Executive Committee member and the Treasurer for ACM SIGMOBILE, the Special Interest Group on Mobility of Systems, Users, Data, and Computing, and is a member of the ACM, IEEE, IEEE Computer Society, IEEE Communications Society, USENIX, and Sigma Xi.