

Exploiting multihop diversity through efficient localized searching with CDMA and route metric-based power control

Jonghyun Kim Stephan Bohacek
kim@eecis.udel.edu bohacek@udel.edu
University of Delaware
Newark, DE 19716

ABSTRACT

Wireless links are subject to large deviations in quality. While such deviations may result in the failure of routes, it also means that some routes may be substantially better than others. In order to utilize these routes, the routing protocol must find these good routes. This paper presents a method for locally searching for routes. Specifically, through passive overhearing of data packets, nodes determine which control packets they should forward. Furthermore, nodes use CDMA with route metric-based power control in order to allow multiple nodes to transmit simultaneously and yet ensure that the control packets advertising the best routes are received. Packet level simulations indicate that these techniques result in significantly better packet delivery ratio, lower delay, and far fewer route failures. For example, in one scenario examined, it was found that the techniques presented resulted in 90% of the connections not requiring any route search after the initial route was found, whereas the other protocols were forced to perform between 4 and 25 route searches for each connection.

1. INTRODUCTION

One of the most important features of wireless networks is the variability of channels. In wireless networking, great pains are taken to mitigate the impacts of the variability of channels. While all layers must cope with the effects of time-varying channels, there has been extensive effort at the network layer. While a common approach is to simply find a new route when the current one fails, there has been extensive effort where precomputed backup paths are found when a route search is performed (e.g., [14] [21] [15] [16] [13] [28] [26] [27] [25]). Thus, when a primary path fails, a new route search is not required.

However, time-varying channels does not only imply that links may break, it also implies that some links are better than others. Indeed, in the context of communication theory, channel diversity means that there may be some chan-

nels between the same transmitter and receiver (but with different antennas) that have better performance than other channels. This diversity is closely related to the stochastic nature of channels. For example, popular models for the channel gain¹ include the lognormal distribution, exponential distribution, and Nakagami. If, through the use of multiple transmit and/or receive antennas, there is a set of channels and these channels can be modeled as independent random variables, then the larger the set of channels, the higher the probability that a good channel can be found. While communication theory provides a clear picture of how to exploit channel diversity (e.g., Chapter 11 of [1]), techniques to exploit diversity in the setting of multihop wireless networks is less well understood.

While there has been little work on exploiting diversity in the multihop scenario, theoretical models of the benefits of diversity in multiple hops has been explored. Specifically, in [4] it was shown that in a five-hop network with lognormal fading and nodes spaced 100 m apart, it is possible to improve performance by more than a factor of 3000. In that work, the performance metric was the SNR over the worst link along the path (i.e., the bottleneck in terms of SNR). In [11], simulations were used to demonstrate performance improvements of several orders of magnitude for several metrics such as throughput, loss probability, power, and energy. To provide some insight into how such large improvements are possible, consider the dynamic range of wireless links such as 802.11 span at least 5 orders of magnitude. Thus, by searching out and utilizing routes composed of good links, it is possible to find good routes, where, like links, the dynamic range of a good route as compared to a bad route may span several orders of magnitude.

While the above efforts have examined theoretical performance improvements offered by diversity, the task of achieving this performance has yet to be fully explored. In [2], a protocol that exploits diversity over a static network was presented. This paper focuses on multihop mobile wireless networks. The multihop case provides important challenges. First, while the prior theoretical work shows that good routes exist, finding these routes requires overhead. Second, as the nodes move, the quality of the links change, and hence a new search must be performed to find the current best route. That is, while good routes may exist, finding and maintaining these good routes may require so much overhead that the performance gains are reduced. Indeed,

¹The received signal power is proportional to transmitted power multiplied by the channel gain.

theoretical analysis has shown that straightforward techniques to exploit multihop diversity require excessive overhead [5]. On the other hand, if a good route is maintained, then route failures will not occur, and hence, a diversity exploiting scheme may yield less route search overhead as compared to a traditional routing approach.

The protocol presented here has two novel components. First, in order to take advantage of diversity, an efficient scheme to maintain good routes is used. Second, an opportunistic method is used to repair routes. Both components have several common characteristics. The route maintenance has the following three key features.

1. CDMA is used to exchange routing information. It should be pointed out that CDMA is difficult to apply in the multihop setting since transmissions are subject to the near-far effect. In particular, CDMA provides fair multiple access only if the transmissions from the different hosts arrive with the same signal strength. While power control makes this possible in the mobile phone setting, it is not possible in MANETs. However, by applying power control that is based on the value of the routing metric that is being advertised by the control packet, the near-far effect can be used to ensure that advertisements of good routes are received, while advertisements of bad routes are lost. This scheme is analyzed in Section 2.3.
2. Instead of searching for the best route in the entire topology, only local route maintenance is performed. Specifically, through passive overhearing of data packets, nodes are able to determine which control packets should be forwarded, and estimate the value of the route metric. Since only the nodes that are nearby the connection are able to overhear data packet transmissions, the route maintenance is localized.
3. Local searches are only performed when the route metric of the current route has been found to drop below a threshold, where the threshold is proportional to the average route metric. Theoretical justification of this technique was presented in [5].

The opportunistic route recovery also makes use of CDMA with route metric-based power control and the local searching, and is discussed in detail in Section 3.

It should be pointed out that there are basic conceptual differences between the approach investigated here and multipath/alternative path routing. One can assume that the initial route used by alternative path schemes is the best route among those found (depending on the protocol, the selection may be more arbitrary). Hence, when the initial route fails, the protocol switches to a route that was decided to not be the best, where the determination of the quality of the route is based on information that may now be stale. We have found that in many cases, this alternative route provides poor quality. The idea here is to not expend overhead searching for alternative routes, but actively and efficiently maintain the active route. A side benefit of the route maintenance is that the same structure used to efficiently maintain routes can be used to repair routes.

2. LOCAL BEST-SELECT PROTOCOL (LBSP)

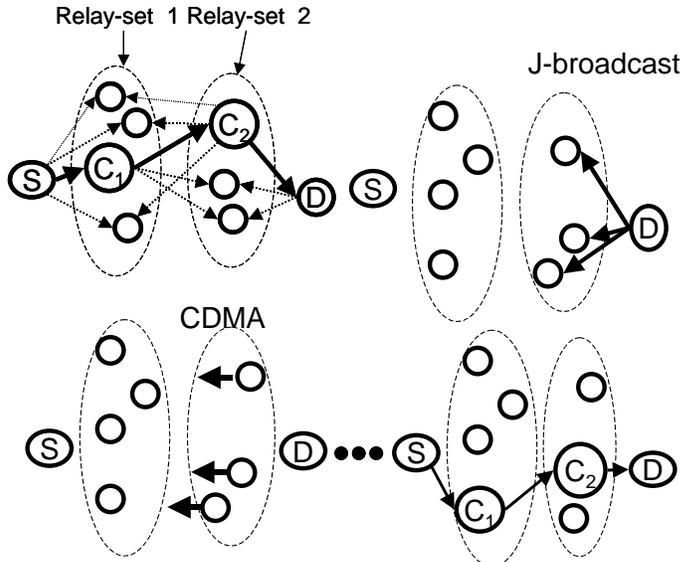


Figure 1: The active path is denoted C_n where n is the number of hops from the source. Nodes that can hear C_{n-1} and C_{n+1} join the n th relay-set. The source is in the 0th relay-set and the destination is in the N th relay-set. When the destination detects that the quality of the route has degraded, it initiates a JBC. The nodes in the $(N - 1)$ th relay-set that receive the JBC packet, update various fields within the packet, and broadcast the packet via CDMA to the next upstream relay-set. This continues until the source receives one or more JBC packets, at which point the route has been updated.

2.1 Overview

In order to demonstrate the utility of exploiting diversity and the several techniques to efficiently exploit diversity, we present a protocol that we refer to as Local Best-Select Protocol (LBSP). This protocol shares many common features with AODV [9], but has highly efficient localized route management. For example, LBSP performs route search in nearly the same way as in AODV. However, the quality of the end-to-end path is continually monitored. If the quality is found to drop below an adaptive threshold, then a local route maintenance is initiated. This route maintenance begins at the destination, and is localized to nodes that are within communication range of the route that is currently being used. A key feature of the route maintenance is that the control packets are broadcast in steps as shown in Figure 1. These broadcasts occur simultaneously, and, in order to ensure that the best routes are found, we use CDMA with route metric weighted power control.

It is important to note that two route metrics are used. First, AODV attempts to minimize the hop count. Second, the local route maintenance seeks to reduce the worst-SNR-to-go, which is explained next. The *worst-SNR-to-go* over a route is the lowest SNR of the all links along the route. We use J_i to denote the worst-SNR-to-go from node i to the destination. The worst-SNR-to-go is a cost-based metric and, in a way that is similar to distance vector routing, obeys

$$J_i = \max_{j \in \text{Neighbors of } i} \min(H_{i,j} - \mathcal{N}, J_j),$$

where $H_{i,j}$ is the channel gain² from node i to node j , and \mathcal{N} is the noise (e.g., thermal noise). Note that all values are in dB.

The worst-SNR-to-go is motivated in several ways. First, it is known that links with high channel gains tend to have long residual lives [6]. Second, if an adaptive bit-rate is used (i.e., the link bit-rate is set according to the link SNR), then the end-to-end throughput is the same as the bit-rate of the slowest link, which is a function of the worst-SNR-to-go from the source to the destination. Third, in many cases, the worst link dominates the end-to-end performance. For example, the end-to-end loss probability is greatly impacted by the link with the worst loss probability, which is related to the worst-SNR-to-go. Similarly, if power control is used, then the total power used is greatly impacted by the power needed to transmit across the link with the worst-SNR. On the other hand, this metric does not account for interference, a topic of future work.

Typically, but not always, LBSP delivers data packets over a route that has been precomputed (as oppose to finding a route on-the-fly). We denote the path over which data packets are delivered as the *active route*. We denote the nodes within the active route as C_n where n indicates the number of hops from the source.

Nodes that are nearby the active route overhear the data transmission and form a local topology that is used for local route maintenance. Specifically, the header of the data packets includes the following information

- the source and destination addresses,
- the number of hops to the source and destination from the transmitting node,

²Due to reciprocity, $H_{i,j} = H_{j,i}$.

- J_{Source} - an estimate of the worst-SNR-to-go from the source to the destination,
- $J_{Transmitter}$ - an estimate of the worst-SNR-to-go from the transmitting node to the destination,
- $J_{Current}$ - the worst SNR experienced by the data packet so far,
- new route bit - which indicates if the source is using a new route,
- CPS , an current estimate of the number of chips per symbol, the use of this information is discussed later.

Note besides the source and destination addresses, which are normally included in data packets, each piece of information within the header only requires a single byte, and hence a total of seven bytes of information is added to each data packet.

Nodes that are able to receive the data transmissions record the source and destination of the packet and also record the number of hops from the source and destination. We refer to the set of nodes that are able to receive transmissions from C_{n-1} and C_{n+1} as the n th relay-set (See Figure 1). Clearly, C_n is also a member of the n th relay-set. However, nodes within the n th relay-set might not be able to communicate with each other. We have found that this feature greatly extends the utility of LBSP, and is in contrast to methods such as the one presented in [3] that requires all members of a relay-set to be within communication range. Finally, a node may be able to hear transmissions from several nodes along the active route. In this case, the node is a member of several relay-sets. Similarly, if a node along the active route is able to hear transmission from nodes that are direct neighbors, these nodes will also be members of multiple relay-sets.

In order to get an idea of the utility of members of the relay-sets, each node estimates its worst-SNR-to-go. Specifically, when node C_{n+1} transmits the data packet, it includes $J_{C_{n+1}}$. When the members of the n th relay set receive this transmission, they measure the SNR and determine their worst-SNR-to-go via $J_i = \min(H_{i,C_{n+1}} - \mathcal{N}, J_{C_{n+1}})$. We note that node C_n also determines J_{C_n} . Hence, the nodes along the currently active route continually maintain the route metric, and hence the J s are reasonably accurate. On the other hand, if the link that is n th hops upstream from the source changes (e.g., the SNR decreases), then it will take $n - m$ data transmissions before the nodes in the m th relay-set update their worst SNR to go. In order to initialize J_{C_n} , the route reply (RREP) packet that initially announces the route includes the worst SNR experienced so far. Hence, the initial values of J_{C_n} are correct.

A key feature of LBSP is that the propagation of route maintenance packets are restricted to nodes nearby the active route. Specifically, only nodes within the relay-sets take part in the route maintenance. The next section discusses route maintenance.

2.2 J-Broadcast and CDMA with Route Metric Weighted Power Control

When the first data packet arrives at the destination, it records $J_{Current}$, the worst-SNR experienced by the packet. We denote this initial value of $J_{Current}$ as $J_{Initial}$. Then, when later data packets arrive, the destination compares

$J_{Current}$ to $J_{Initial}$. If $J_{Current} < J_{Initial} - J_{Threshold}$, then the quality of the current path is questionable, and hence a route maintenance is performed. Here $J_{Threshold}$ is a design parameter that is discussed later. The route maintenance is initiated by the destination and, essentially, propagates control messages to the source. The packets are restricted to the nodes within the relay-sets for the source-destination pair for which the route is being maintained. As these packets are propagated toward the source, the route that provides the highest worst-SNR-to-go from the source to destination is found. A key aspect of this route maintenance is that the control packets are transmitted efficiently with CDMA and power control. We provide details of this scheme next.

We refer to route maintenance control packets as J-broadcast (JBC) packets. The route maintenance is initiated when the destination broadcasts a JBC packet. The JBC packet contains the following information,

- source and the destination,
- the number of hops from the source,
- the value of $J_{Transmitter}$ (which, for the destination is always 0 dB),
- $J_{Current}$, the worst-SNR-to-go from the source to the destination,
- MAX_NUM_JBCS, which contains the maximum number of JBC packets received by a node,
- and the transmit power (which, in the case of the destination, is the full transmit power).

For the sake of discussion, suppose that the destination is N hops from the source. Thus, the destination broadcast the JBC packet to the members of the $(N - 1)$ th relay-set receives this JBC packet. Upon receiving the JBC packet, these nodes measure the SNR and update their worst-SNR-to-go, which, since the nodes are exactly one hop from the destination, is the same as the SNR of the received JBC packet. We denote node i 's worst-SNR-to-go as J_i .

If node i within the $(N - 1)$ th relay-set that has $J_i \leq J_{Current}$, then the JBC packet is dropped. This technique of dropping of JBCs when J is small is referred to as *J-threshold test*. The impact of this technique is examined below.

If a node has $J_i > J_{Current}$, then a JBC is prepared. This packet contains the same information as the JBC that was received from the destination, but with an updated value of $J_{Transmitter} := J_i$. Furthermore, we include a data field called the MAX_NUM_JBCS that is set to one, indicating that the JBC packet was received from a single transmitter. This field is later used to adjust the chips per symbol in the CDMA transmissions. Adjusting the number of chips per symbol is a second techniques used in the efficient propagation of JBC packets and is also examined below.

After the eligible nodes prepare the JBC packets, they are broadcast using CDMA with the number of chip per symbol equal to the value that was provided in the data packet and the transmission power is set according to

$$TransmitPower_i = \min(MAX_POWER, J_i - J_{Current} + TARGET_POWER), \quad (1)$$

where MAX_POWER and $TARGET_POWER$ are design parameters that are analyzed shortly and where we assume that all values are in dBm. Note that the transmit power depends on the value of J_i , which is the value of the route metric. This route metric-based power control is the third technique used to efficiently propagate JBC packets, and is also further discussed below.

The nodes within the $(N - 2)$ th relay-set then receive the transmissions from the $(N - 1)$ th relay-set. Of course, these nodes receive the JBC packets at various signal strengths. Specifically, node j receives the transmission from node i with signal strength $H_{i,j} + TransmitPower_i$, where $H_{i,j}$ is the channel gain (in dB) from node i to node j . Note that the probability of packet error via CDMA depends on the relative received signal strengths. We denote the relationship between the vector of received signal strengths and the probability of transmission error from the i th source to be $E(i, \vec{P}, CPS)$, where \vec{P} denotes the vector of received signal strength, and CPS denotes the number of chips per symbol (See Appendix E in [19] for details on E). A sample of E is shown in Figure 2.

Upon receiving and decoding the JBC packets, the nodes in the $(N - 2)$ th relay-set update their worst-SNR-to-go and record the next hop along the path with the largest worst-SNR-to-go. Specifically, node j updates J_j according to

$$J_j = \max_{\{i | \text{the JBC from node } i \text{ was decoded correctly}\}} \min(H_{i,j} - \mathcal{N}, J_i), \quad (2)$$

where the node j 's best next hop is the node i that achieves the maximization.

These nodes then repeat the procedure performed by the nodes in the $(N - 1)$ th relay-set, i.e., perform the J-threshold test, perhaps form a JBC packet with $J_{Transmitter}$ and the MAX_NUM_JBCS fields updated, and then transmit with power given by (1). This process repeats until the source receives the JBCs, at which point the source determines its best next hop, and hence the route is updated. If the source does not receive a JBC packet, the old route is continued to be used.

2.3 JBC Propagation Performance

2.3.1 Overview

While CDMA has the ability to receive data from multiple sources simultaneously, as illustrated in Figure 2, if the received powers of the transmissions are greatly different and/or the number of chips per symbol (CPS) is too small, then transmission errors occur with high probability. This problem with CDMA is also present in mobile phones, however, in that setting it is possible to assign transmission powers such that each transmission is received with the same signal strength. On the other hand, in general, it is not possible to adjust the transmit powers from multiple transmitters so that they are received with equal powers at *multiple* receivers. Nonetheless, power control, specifically, (1), can be effectively used to support the goal of communicating the best route.

In order to provide motivation of (1), note that if either $H_{i,j}$ is small or $TransmitPower_i$ is small, then the JBC from node i is unlikely to be received by node j . On the other hand, if $H_{i,j}$ is small or $TransmitPower_i$ is small, then it is unlikely that the best route includes the link $\vec{i, j}$.

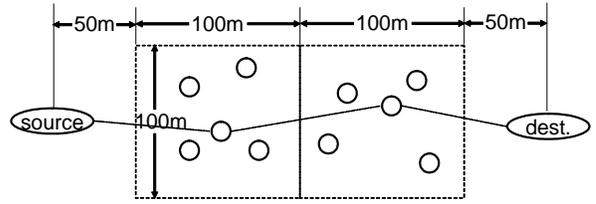


Figure 3: The Topology for Analyzing the Performance of CDMA with Route Metric-Based Power Control. Relay-sets are made up of nodes uniformly distributed over $100\text{m} \times 100\text{m}$ squares. The first relay-set is centered 100 m from the source, while the second relay-set is 200 m from the source and 100 m from the destination. While the above shows five nodes in each relay-set, in the analysis, different numbers of nodes are considered.

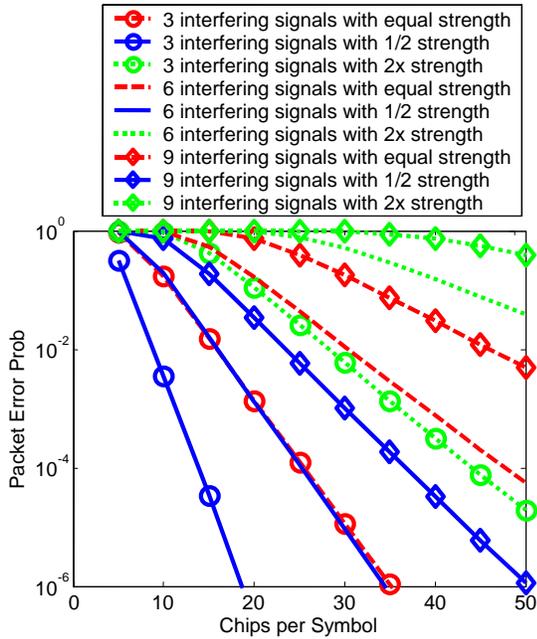


Figure 2: The error probability of CDMA depends on the relative signal strengths of the interfering signals and the number of chips per bit. Above shows the packet error probability (40 bytes per packet) when there are 3, 6, and 9 interfering signals and when the signal under test has the same strength, half the strength, or twice the signal strength as each of the interfering signals.

More precisely, if $H_{i,j}$ is small, then $\min(H_{i,j} - \mathcal{N}, J_i)$ is small, and if the $TransmitPower_i$ is small, then J_i is small, and hence $\min(H_{i,j} - \mathcal{N}, J_i)$ is small. In either case, node i will likely not achieve the maximum in (2), and hence node j 's next hop is not node i . Therefore, if a JBC is received with a small signal strength, then the inability to correctly decode the transmission will not have an impact on the ability of the source to find the best path. Furthermore, by reducing the transmit power of nodes with a low worst-SNR-to-go results in less interference for the transmissions of JBCs advertising high worst-SNR-to-go.

Unfortunately, while there is significant correlation between received signal strength and the worst-SNR-to-go, they are not perfectly correlated. Specifically, it is possible that a JBC may be received with a moderately high signal strength, and yet may advertise a route with moderately low worst-SNR-to-go. To see this, note that the JBC from node i advertises the route with worst-SNR-to-go equal to $\min(H_{i,j} - \mathcal{N}, J_i)$, whereas the received SNR is *proportional* to $H_{i,j} + J_j - \mathcal{N}$. This difference between the addition and minimization operations results in some JBCs with moderately good worst-SNR-to-go to be received with larger signal strength than the JBC advertising the largest worst-SNR-to-go. Hence, LBSP must maintain a large enough number of chips per symbol so that JBCs that advertise very good routes are properly decoded in the presence of the interference due to JBCs advertising moderately good routes. Next we present numerical results to further investigate the performance of CDMA for delivering routing control messages.

2.3.2 Numerical Analysis of JBC with CDMA and Route Metric-Based Power Control

In order to analyze the behavior of CDMA we consider the idealized topology shown in Figure 3 (more realistic topologies are considered in Section 4). The channel gain (in dB) between nodes that are d meters apart is modeled as $2.3 \times 10 \log_{10}(d) + X$, where X is a Gaussian random variable with mean 0 and standard deviation of 12 dB (see [19]).

In general, the performance of CDMA is examined as follows. It is assumed that all nodes in the second relay-set decode the JBC from the destination. These nodes perform the J-threshold test, that is, they compare their worst-SNR-to-go (i.e., the SNR to the destination) to $J_{Current}$,

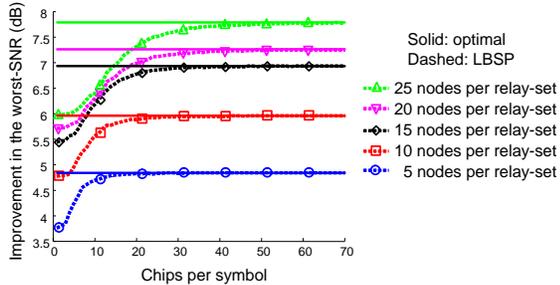


Figure 4: Improvement in the worst SNR (i.e., the SNR of the link with the lowest SNR of all links along the route). The solid lines show the average improvement offered by diversity over the geographical shortest path. The dashed lines show the performance achieved when the CDMA-based scheme is used to exchange JBC routing control packets. The different sets of curves are for different numbers of nodes in each relay-set (see Figure 3 for details).

which is set to the worst-SNR-to-go from the source to the destination via the geographically shortest path. Next, for the nodes that pass the J-threshold test, the JBC transmit power is calculated via (1). The probability that a JBC from an eligible node in the second relay-set is correctly decoded by a node in the first relay-set depends on the vector of received signal strengths and is provided in Appendix E of [19]. Based on a realization with these probabilities, the set of JBCs that are correctly decoded by each node in the first relay-set is simulated. For each node in the first relay-set, the worst-SNR-to-go is calculated, and the steps performed by the nodes in second relay-set are repeated. At this point, the worst-SNR-to-go from the source to the destination is calculated and compared to the worst-SNR-to-go over the geographically shortest path. More specifically, the improvement of the worst-SNR-to-go is recorded. If the source does not correctly decode any messages, then there is no improvement over the geographically shortest path. All figures below show the average performance improvement in dB.

Figure 4 shows the impact of the number of chips per symbol. This figure shows marked curves and straight lines. The straight lines indicate the improvement offered by diversity, that is, the improvement when the best path is used as oppose to the path that minimizes the geographical distance. As expected, as the number of nodes in each relay-set increases, the improvement offered by diversity increases. This improvement does not vary with the number of chips per symbol; it represents the upper limit in performance improvement assuming global topology knowledge.

The curves in Figure 4 show the average performance improvement when control packets are transmitted with the CDMA-based scheme discussed above. As expected, as the number of CPS increases, the control packets are decoded with higher probability and hence a better quality route is found. Furthermore, as the number of nodes increases, a larger number of CPS is required to achieve, or nearly achieve, the optimal performance improvement offered by diversity.

There are three mechanisms that are used in the broad-

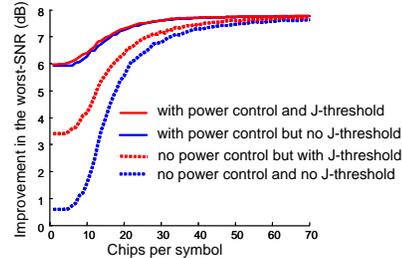


Figure 5: A comparison of schemes to improve the performance of transmission of routing control packets with CDMA. This example shows the performance when there are 25 nodes in each relay-set and there are three hops. The upper most curve is the same as the one shown in Figure 4.

cast of JBCs. First, CDMA is used with several chips per symbol. Second, power control is used to attempt to allow the reception of JBCs that advertise the best routes to be decoded correctly. And third, the J-threshold test is applied to reduce the number of broadcasts advertising routes with low worst-SNR-to-go. Figure 5 examines the impact of these three techniques when there are exactly 25 nodes in each relay-set. It can be seen that if the number of CPS is large enough, then the optimal performance is achieved without making use of the other two mechanisms. However, if the J-threshold test is used, then the effectiveness of the CDMA transmissions is significantly improved. The reason for the improvement is that the J-threshold technique stops nodes with low worst-SNR-to-go from broadcasting, and hence reduces the interference experienced by JBCs advertising large worst-SNR-to-go.

Figure 5 also shows the performance improvement when power control is used, but without the J-threshold test. We see that as compared to the J-threshold test alone, power control alone significantly improves the performance. Interestingly, if power control is combined with the J-threshold test, no further improvement is achieved. Considering this last observation (which holds for relay-set sizes from 5 to 100), there seems to be little justification for using the J-threshold test. However, since this technique does reduce the number of transmissions, it reduces the overhead and saves power. Hence, LBSP uses both route metric-based power control and the J-threshold test.

Figures 4 and 5 clearly show that the performance improvement depends on the relay-set size and the number of CPS. Thus, in order to correctly apply this technique, the correct number of CPS must be determined. Here we examine the relationship between performance, the relay-set size, and the number of CPS. Specifically, the right-hand frame in Figure 6 shows the performance improvement with CDMA, power control, and J-threshold test (this is essentially the same as Figure 4). On this figure are marks indicating the number of CPS that achieve 90%, 99% and 99.7% of the optimal performance improvement offered by diversity. In Figure 7, the number of CPS required to achieve 99% of the optimal performance is plotted as a function of the number of nodes in the relay-sets. This relationship is surprisingly

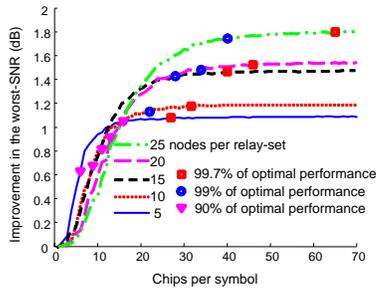


Figure 6: Performance vs. the Number of Chips per Symbol (CPS) and the size of each relay-set. The curves show the relationship between the performance and the number of CPS. The marks indicate where a specific fraction of the optimal performance is achieved.

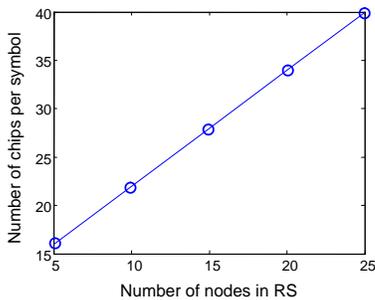


Figure 7: The relationship between the number of nodes in each relay-set and the number of chips per symbol required to achieve 99% of the optimal performance.

linear, specifically, it is closely approximated by

$$\text{Number of chips per bit} = 1.2 \times \text{Relay-set size} + 10.$$

Based on this analysis, the above relationship is used by the source to determine the number of chips per bit. Specifically, when a JBC is received by the source, it estimates the maximum number of nodes that broadcasts JBCs from the `MAX_NUM_JBCS` field. The CPS that the source selects is included in the data packet header so all nodes within the relay-sets are aware of the number of CPS.

There is one final aspect of the power control that remains. Specifically, (1) indicates that the transmit power is bounded by the constant, `MAX_POWER`. In practice, transmit power is always bounded. Hence, the impact of this constant must be examined. Note that if the best route has a worst-SNR-to-go that is approximately the same as $J_{Current}$, then the transmission power of the JBC that advertises this route is `TARGET_POWER` dBm. Of interest is the difference

$$\Delta P := \text{MAX_POWER} - \text{TARGET_POWER}.$$

Note that if $\Delta P = 0$, then only the J-threshold test is used. To see this, note that nodes do not transmit if $J < J_{Current}$. And hence, never transmit below `TARGET_POWER`. On

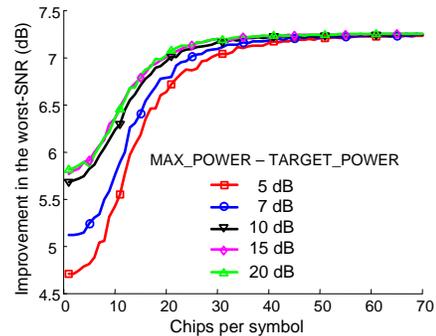


Figure 8: The Impact of Limiting the Maximum Transmit Power. This example considers 20 nodes in each relay-set. The difference between the maximum transmit power and the target power ranges from 5 to 20 dB.

the other hand, if $\text{MAX_POWER} = \text{TARGET_POWER}$, then nodes never transmit with power above `TARGET_POWER`. Thus, the larger the difference ΔP , the more of an impact the power control may have.

Figure 8 shows the impact of ΔP for 20 nodes in each relay-set. It can be seen that if the difference is beyond 15 dB, there is little further improvement. A similar relationship holds for other sized relay-sets. Hence, LBSP uses $\text{TARGET_POWER} = \text{MAX_POWER} - 15$ dB, where `MAX_POWER` is the maximum allowable transmission power.

3. OPPORTUNISTIC ROUTE RECOVERY

If data packets are sent frequently, then the route metric is frequently reestimated, and route maintenance is effective in keeping the worst-SNR-to-go large. However, if packets are sent infrequently, then the route may break before the weakening link can be detected. However, the relay-sets and CDMA can be used for opportunistic route repair as follows.

Suppose that C_n , the node along the active route that is n hops from the source, attempts to transmit a packet over its best next hop, but the transmission fails. Upon detecting the failure, C_n will broadcast a request to send (RTS). This RTS will include the addresses of the source and the destination of the connection and also specifies that the members of the $(n + 1)$ th relay-set are the desired receivers. Upon receiving the RTS, the members of the $(n + 1)$ th relay-set check if the channel is idle and, if so, transmit a clear to send (CTS). However, these transmissions use CDMA with the power control described above. The CTS also includes the transmitters' current estimate of the worst-SNR-to-go (see Section 2 for discussion on how the nodes maintain this estimate). Upon receiving the CTSs, C_n selects the node that provides the largest worst-SNR-to-go as its best next hop, and transmits the data packet to this node.

If after detecting a failure broadcasting the RTS, no CTSs are received and if C_n is not a member of another relay-set (besides the n th relay-set), then a route error (RERR) packet is generated and delivered to the source, in the same way as is done in AODV. On the other hand, if the node is a member of another relay-set, then further recovery steps can be taken. Specifically, if C_n is a member of the $(n + k)$ th

relay-set for $k \geq -1$, then the node can broadcast a RTS and specify the receivers to be members of the $n+k+1$ relay-set. This rebroadcast of RTSs can occur repeatedly until a CTS is received, or until all relay-sets that C_n is adjacent to have been tried, at which point, a RERR message is generated.

It should be pointed out that C_n will not send an RTS to any relay-set that is upstream, i.e., relay-set $(n-k)$ for $k > 0$. It may, however, attempt to send an RTS to members of the n th relay-set, which C_n is also a member of. On the other hand, this is action cannot be taken at the source or destination. That is, following the definition of the relay-set above, the source is the only member of the 0th relay-set. Similarly, if there are N hops from source to destination, then the destination is the only member of the N th relay-set. This limited flexibility results in the first or the last hop being more susceptible to link failure. To mitigate this, we allow nodes to join to 0th relay-set if they can hear C_1 and hear the source. Similarly, a node can join the N th relay-set if they can hear C_{N-1} and the destination. Thus, if the source is unable to send a packet to C_1 , and is unable to reach any member of the first relay-set, it attempts to reach the members of the 0th relay-set. C_{N-1} behaves similarly.

If a node performs a successful route recovery (i.e., it finds an alternative best next hop), then the best next hop is selected for further transmissions. Note that in this case a node may be acting like C_n and C_{n+k} for $k \geq -1$. For example, C_n may receive the data from C_{n-1} , but then have a best next hop in relay-set n , or it may have a best next hop in the $n+2$ relay-set. While this does not have any adverse impact on data delivery, it does confuse which nodes join the relay-sets. For example, if a node receives data from C_{n-1} and transmits data to C_{n+2} , then a is able to only hear the transmissions from this node and C_{n+2} will join the $(n+1)$ th relay-set. Similarly, if a node hears transmissions from this node and from C_{n-1} , it will join the n th relay-set. Figure 9 illustrates this behavior.

It should be pointed out that typically after a route recovery is performed, the SNR experience by the packet degrades enough that a JBC results. And hence a new route is found without a node acting as multiple hops along the path. On the other hand, if the source or C_{N-1} (the last hop before the destination) perform a route recovery, then the length of the path is increased.

While the above schemes may extend the length of the route, it is sometimes useful to shorten the route. To this end, if a node along the active route, say C_n , observes that it the SNR from transmissions from the downstream node, C_{n+2} surpass a threshold, then the data is directly transmitted to C_{n+2} , bypassing C_{n+1} . We selected 35 dB channel loss as the threshold.

In Section 4.3 we compare the performance with and without these additional route recovery features. As will be seen, these techniques have a significant impact on the performance.

4. PACKET-LEVEL SIMULATIONS

4.1 Scenarios Examined

We examine the performance of the above scheme in an indoor urban scenario. Specifically, we consider communication within a building. This scenario is of interest in the urban mesh networks that are currently being deployed throughout the world (e.g., [17], [10], [24], and [8]). How-

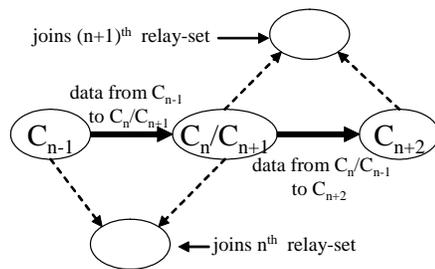


Figure 9: Route recovery may result in a node being in the active route, but acting as though it is in two or more hops along the route. The above illustrates the case when a node act as C_n and C_{n+1} . Due to information in the data and ACKs, nearby nodes are able to detect when this has occurred. In this case, the rule of joining a relay-set are maintained, i.e., if a node can hear a transmission from C_{k-1} and C_{k+1} , then the node joins the k th relay-set. The fact that one of the nodes it overhears is actually C_n and C_{n+1} (as shown above) increases the opportunity to join a relay-set.

ever, due to high cost of installing base stations in every building in the city, these networks do not provide coverage into buildings. For example, in Philadelphia and San Francisco, the network is planned to only reach at least one part of 90% of the buildings in the cities. In order to achieve ubiquitous wireless access both indoors and outdoors, connections must be relayed by the mobile hosts³. Thus, indoors is a realistic deployment scenario of MANET routing. Furthermore, while there has been extensive investigation of open-space propagation and random-way point mobility, relatively little attention has been focused on indoor multihop wireless communication. Hence, the following will also provide insight into the performance of other MANET routing protocols indoors.

The specific scenarios examined is a ten story building with $200\text{m} \times 100\text{m}$ footprint. The node densities are shown in Table 1. We consider office building mobility. Specifically, following the studies of office worker time-use [18], [20], an office worker mobility model was developed. More details on the mobility model can be found in [12] and [7].

To examine the performance under more dynamic mobility, a vehicle network was considered. Specifically, a $2\text{km} \times 2\text{km}$ region of downtown Chicago was modeled. A random walk mobility was assumed, where vehicles turn with probability 0.1. Furthermore, once a vehicle exits the city, it enters the city according to a Poisson process. More details of the vehicle model can be found in [12] and [7].

Propagation is an important aspect of MANET performance. The attenuation factor model has been shown to be a realistic model of indoor propagation (See [12], [23], [7] for more details). One important parameter of the ARF model is attenuation due to propagation through a floor. Following the findings presented [19], we assume that there is an added

³Another possibility is that the when a mobile node moves indoors, it would switch to an indoor network provided by the landlord, for example. However, this approach is difficult since users must have agreements with each landlord where they might roam.

Table 1: Indoor Scenarios

Scenario 1	64 nodes
Scenario 2	128 nodes
Scenario 3	256 nodes
Scenario 4	512 nodes
Scenario 5	1024 nodes

Table 2: Outdoor Scenarios

Scenario 1	64 nodes
Scenario 2	128 nodes
Scenario 3	256 nodes

attenuation of 35 dB when crossing a single floor and 45 dB when crossing four or more floors. Since we set the floors to be 3.5 m apart, propagating 4 floors results in 23 dB loss due to the distance and resulting in a total of 68 dB channel loss. On the other hand, typical 802.11a/b/g transmitters and receivers (these simulations are based on 802.11b at 2Mbps along with the modifications described above) are only able to tolerate 65 dB of attenuation. Hence, it is not possible to propagate four floors or more, and propagating three floors is only possible if the transmitter and receiver are nearly vertically aligned⁴. All mobility and propagation traces used for these simulations are available [7].

For the vehicles, the propagation was modeled using ray-tracing. Hence, the impact of buildings was included. See [23], [7] for more details on this model.

This investigation examines constant bit-rate (CBR) traffic. Specifically, we focus on sending 1KB packets every 50ms or 400 ms (high data rate and low data rate, respectively). All connections are tested for 5 minutes. This connection duration provides insight into voice-over-IP applications and multimedia streaming.

In order to gain insight into the performance, the performance of LBSP is compared to AODV [9] and AOMDV [15]. AOMDV is a multipath extension to AODV. During route search, nodes may accept and forward several route reply (RREP) messages, as oppose to AODV, where a node may only forward one RREP. Nodes also maintain several routes to a destination. One point examined by this paper is whether better performance is achieved from precomputing alternative paths, such as is done in AOMDV, or by maintaining paths, as is done in LBSP.

4.2 Simulation Results

We begin by examining the packet delivery ratio (PDR). Figure 10 and 11 shows the PDR for the three protocols and the node densities shown in Tables 1 and 2. As expected, AOMDV provides better a higher PDR than AODV, however, LBSP provides significantly better PDR than AOMDV. Also, as expected, in the outdoor mobility case, LBSP suffers from a lower PDR than the indoor case. However, in both cases, the packet loss probability is below 0.25%, which is approximately an order of magnitude less than the other protocols. As a point of reference, if the loss probability of

⁴The propagation between floors can be greatly enhanced by nearby buildings. For example, the signal may propagate outside, reflect off of a building, and then transmit back into the building. While such a possibility is included in the UdelModels [7], it was not considered here, where it was assumed that the building stood in isolation.

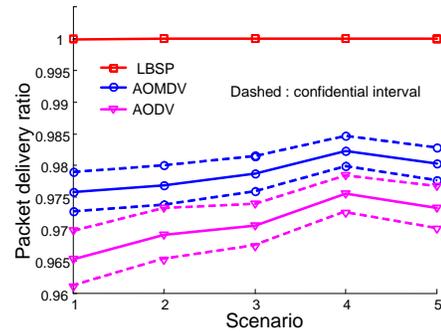


Figure 10: Packet Delivery Ratio for LBSP, AODV, and AOMDV and for different node densities in the indoor MANET. The 80% confidence interval is also shown. In the case of LBSP, nearly every packet was successfully delivered, and hence the confidence interval is very small and indistinguishable from the mean.

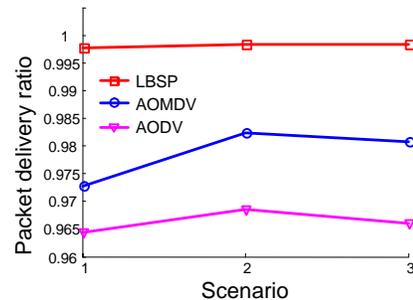


Figure 11: Packet delivery ratio for vehicles.

exceeds 2% TCP begins to suffer from timeouts results in greatly reduced throughput [22].

Route failure is an important cause of packet delivery failure. Figures 12 and 13 shows the average number of route searches for each connections the indoor and vehicle networks respectively. Considering the indoor case, we see that while the connection duration is only 300 sec., AODV requires approximately 15 route searches. That is, a new route must be found once every 20 sec. AOMDV performs far better than AODV with 4.2 route searches per connection. However, LBSP only needs 1.09 route searches per connection. Note that this includes the initial route search. Hence, with LBSP, over 90% of the connections suffered no route failures. As expected, the vehicle network suffers from more route failures than the indoor network. However, LBSP has 5 times fewer failures than AOMDV, and approximately 10 times less than AODV.

Another view of the number of route searches is shown in Figure 14, which shows the overhead due to route request, route reply, and route error messages. This figure focuses on the indoor MANET, the vehicle MANET is similar. This figure shows that AODV suffers from an order of magnitude more overhead than LBSP and AOMDV suffers from approximately 5 times more overhead than LBSP.

Another important aspect of overhead is the JBC packets. In order to examine the amount of overhead, we define

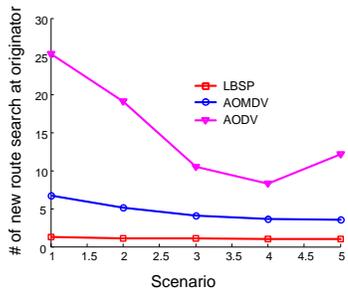


Figure 12: Average number of route searches.

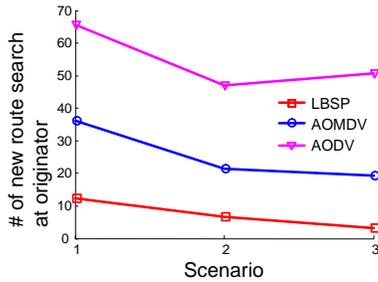


Figure 13: Number of route failures for the vehicle network.

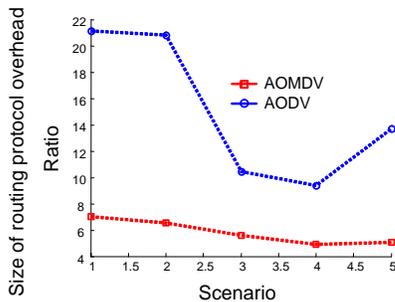


Figure 14: Ratio of the number routing overhead packets sent by AODV and AOMDV to the number sent by LBSP.

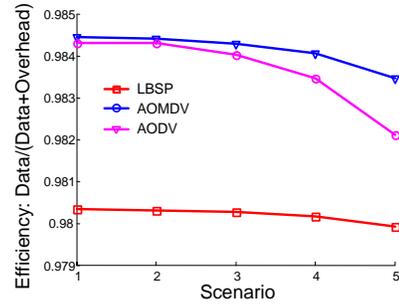


Figure 15: Efficiency of the routing protocols for the indoor MANET. The efficiency is defined to be the total number of data bytes sent by any node divided by the total number of any type of bytes sent by any node.

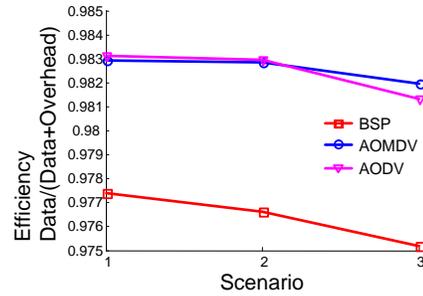


Figure 16: Efficiency of the routing protocols for the vehicle MANET.

efficiency to be the total bytes of data transmitted by any node divided by the number of any type of bytes transmitted. Figures 15 and 16 show the efficiency for the indoor and vehicle MANETs, respectively. It can be seen that LBSP has a lower efficiency than the other protocols. However, note that the efficiency for all protocols is very high, hence nearly all packets transmitted are data packets, and hence little bandwidth or power is consumed with overhead packets. Furthermore, these figures is based on counting bytes. However, when CDMA is used, the channel is occupied by several transmissions at the same time. Furthermore, the power control is based on the assumption that some of these packets will not be decoded by any destination. Hence, the duration that channel is occupied less than is shown.

Another important performance metric is the delay experienced by packets. Figure 17 shows the delay for the different protocols in the indoor MANET. Packet delay is affected by two things, namely, route searches, which delay the delivery until a route is found, and the length of the route. In the case of AODV, the packet delivery is mostly impacted by the route searches, and hence the variation in delay as a function of node density is similar to the what is displayed in Figure 12. On the other hand, LBSP suffers from very few route failures, but as node density increases, a shorter path can be found, and hence the packet delay is slightly decreased. AOMDV performs between LBSP and AODV.

LBSP use the worst-SNR-to-go as a route metric. Hence,

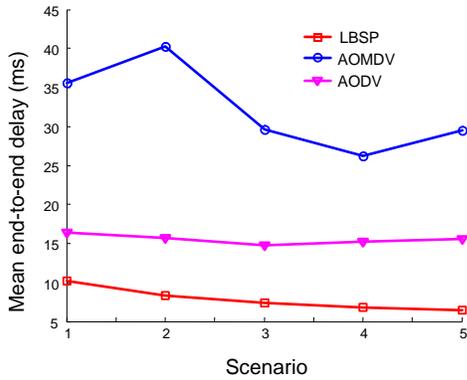


Figure 17: Packet delay for different scenarios in the indoor MANET.

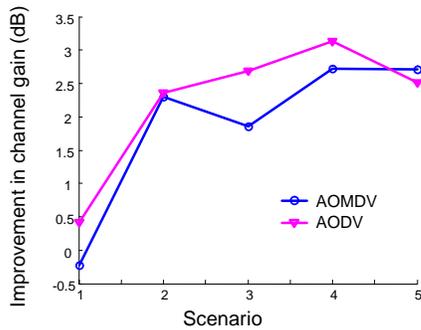


Figure 18: The worst SNR of a route is the SNR of the link along the route. For each protocol, the average value of the worst SNR was computed for each packet delivered. The above plot shows the ratio of these worst SNRs.

one would expect that LBSP uses better routes than AODV and AOMDV. Figure 18 shows the average value of the worst-SNR-to-go from the source to destination, where the average is taken over each packet delivered and each connection in the indoor MANET. The figure shows that LBSP utilizes routes that are between 50% to 100% better. While this difference is large, it is quite a bit less than the achievable differences offered by diversity. The reason for the limited improvement is that LBSP is conservative in exploring routes. Specifically, routes are only explored when needed, i.e., when the worst-SNR-to-go from the source to destination drops below the SNR found during the last JBC. While this approach reduces the gains in SNR, it successfully achieves high packet delivery ratio and reduces route searches, and also minimizes the overhead associated with more fully exploiting diversity.

4.3 Performance with Different Protocol Parameters

While Section 2.3 examined the protocol parameter selection related to the CDMA and power control, there are two aspects that remain to be investigated. Specifically, the impact of the route recovery techniques described in Section 2.2, and the value of $J_{Threshold}$ use to trigger a JBC. First,

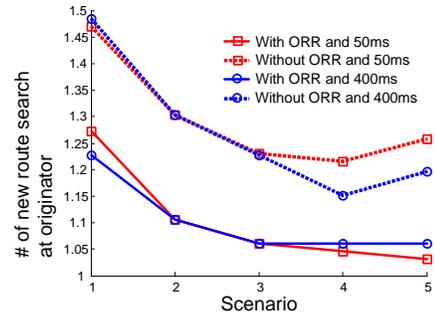


Figure 19: Performance of LBSP with and without route repair. Packets were every 50 msec and every 400 msec.

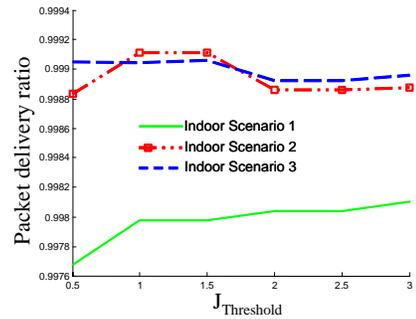


Figure 20: The impact of $J_{Threshold}$. Here the indoor MANET is considered for scenarios 1, 2 and 3.

route recovery is examined.

Figure 19 shows the average number of route searches for each connection with and without the route recovery scheme. Furthermore, we consider packets send every 50 msec and every 400 msec. At the lower data rate the route is probed less often, and hence the route maintenance is less effective. Recall that each connection requires an initial route search. Hence, the route recovery techniques reduce the number of extra route searches by a significant amount. For example, in the case of high node densities, only 3% of connections suffer from a route failure if route recovery is used, but approximately 30% of connections suffer a route failure if route recovery is not used.

As described in Section 2.2. when the SNR experienced by a packet drops significantly below the SNR experienced by the first packet sent after the route has been found or after a JBC has been completed, then a JBC is initiated. Specifically, if $J_{Current} < J_{Initial} - J_{Threshold}$, then JBC is performed. Figure 20 shows the performance for various values of $J_{Threshold}$. These simulations were performed without the route recovery techniques discussed in Section 3.

Surprisingly, we see little difference in the performance. The reason for this is that small variation in performance is that links between stationary nodes are stable, but break soon after nodes move. The node movement can be detected via a change in the SNR. Thus, the link failure is predicted by a change, the size of the change does not provide significant additional evidence of the impending link breakage.

5. CONCLUSIONS

A MANET routing protocol that used signal strength to maintain route has been presented. This protocol has several unique features. In order to reduce overhead, it used CDMA with route metric-based power control to exchange routing control packets. Extensive analysis of the performance of such an approach shows that the CDMA can be effectively used for the exchange of control packets. A second unique feature of the protocol is that it uses passive overhearing of data packets to learn about nodes that may be potential relays. A third technique used by this protocol is that if the performance of the current route is found to degrade, then a localized and efficient route search is performed. Packet simulations show that in terms of packet delivery ratio, delay, and the number of route failures, these techniques result in performance that is better than AODV and AOMDV, a multipath version of AODV. On the other hand, the presented technique does result in more overhead. However, the overhead of all protocols nearly the same, and represents only a small fraction of the bytes transmitted.

6. REFERENCES

- [1] BARRY, J. R., LEE, E. A., AND MESSERSCHMITT, D. G. *Digital Communication*. Kluwer Academic Publishers, Boston, 2004.
- [2] BISWAS, S., AND MORRIS, R. ExOR: Opportunistic multi-hop routing for wireless networks. In *Sigcomm05* (2005).
- [3] BLETSAS, A., KHISTI, A., REED, D. P., AND LIPPMAN, A. A simple cooperative diversity method based on network path selection. *IEEE Journal on Selected Areas of Communication* (2005).
- [4] BOHACEK, S. Achievable performance improvements provided by cooperative diversity. In *4th Intl. Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)* (2006).
- [5] BOHACEK, S., BLUM, R., CIMINI, L., GREENSTEIN, L., AND HAIMOVICH, A. The impact of the timeliness of information on the performance of multihop best-select. In *Military Communications Conference (Milcom)* (2005).
- [6] BOHACEK, S., ILIC, A., AND SRIDHARA, V. On the predictability of link lifetimes in urban MANETs. In *3rd Intl. Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOPT)* (2005).
- [7] BOHACEK, S., SRIDHARA, V., AND KIM, J. UDelModels. <http://udelmodels.eecis.udel.edu>.
- [8] CASS, S. Viva mesh vegas. *IEEE Spectrum* (2005).
- [9] CHARLES E. PERKINS, E. M. R. Ad hoc on-demand distance vector routing. *2nd IEEE Workshop on Mobile Computing Systems and Applications* (1999), 90–100.
- [10] CITY, AND OF SAN FRANCISCO, C. Request for proposals 2005-19 - TechConnect: Community wireless broadband network. Available at http://www.sfgov.org/site/tech_connect_index.asp, 2005.
- [11] KIM, J., AND BOHACEK, S. Selection metrics for multihop cooperative relaying. In *Proceeding of the Fourth Annual Mediterranean Ad Hoc Networking Workshop (MedHoc)* (2005).
- [12] KIM, J., SRIDHARA, V., AND BOHACEK, S. Realistic simulation of urban mesh networks - part i: Urban mobility. *Submitted* (2006).
- [13] LEE, S.-J., AND GERLA, M. AODV-BR: backup routing in ad hoc networks. In *IEEE WCNC* (2000), pp. 1311–1316.
- [14] LEE, S.-J., AND GERLA, M. Split multipath routing with maximally disjoint paths in ad hoc networks. In *ICC* (2001).
- [15] MARINA, M. K., AND DAS, S. R. Ad hoc on-demand multipath distance vector routing. Tech. rep., SUNY - Stony Brook, 2003.
- [16] NASIPURI, A., AND DAS, S. R. On-demand multipath routing for mobile ad hoc networks. In *Proceedings of IEEE International Conference on Computer Communications and Networks ICCCN* (1999), pp. 64–70.
- [17] NEFF, D. Wireless philadelphia town hall meeting, November 2004. available at <http://207.245.67.199/philagotit/wifilive2.wmv>.
- [18] PANKO, R. R., AND KINNEY, S. T. Meeting profiles: Size, duration, and location. In *Proceedings of The 28th Annual Hawaii International Conference on Systems Science* (1995).
- [19] RAPPAPORT, T. *Wireless Communication*. Pearson Education, 2002.
- [20] ROMANO, N. C., AND NUMAMAKER, J. F. Meeting analysis: Findings from research and practice. In *Proceedings of Teh 34th Hawaii International Conference on Systems Science* (2001).
- [21] SAMBASIVAM, P., MURTHY, A., AND BELDING-ROYER, E. M. Dynamically adaptive multipath routing based on AODV. In *MedHocNet* (2004).
- [22] SHAH, K., AND BOHACEK, S. TCP throughput and timeout – steady state and time-varying dynamics. In *IEEE Global Telecommunications Conference (GLOBECOM)*, (2004).
- [23] SRIDHARA, V., AND BOHACEK, S. Realistic simulation of urban mesh networks - part II: Urban propagation. *Submitted* (2006). <http://udelmodels.eecis.udel.edu>.
- [24] TIMES, T. Ma says taipei to be wireless in two days, December 19, 2005.
- [25] TSIRIGOS, A., AND HASS, Z. J. Multipath routing in the presence of frequent topological changes. *IEEE Communications Magazine* (2001), 132–138.
- [26] VUTUKURY, S., AND GARCIA-LUNA-ACEVES, J. J. MDVA: a distance-vector multipath routing protocol. In *Proceedings of IEEE INFOCOM'01* (2001).
- [27] WANG, L. Multipath source routing in wireless ad hoc networks. In *Canadian Conference on Electrical and Computer Engineering* (2000), pp. 479–483.
- [28] ZHANG, L., ZHAO, Z., SHU, Y., WANG, L., AND YANG, O. W. Load balancing of multipath source routing in ad hoc networks. In *Proceedings of IEEE ICC'02* (2002).