

A Comparison of Opportunistic and Deterministic Forwarding in Mobile Multihop Wireless Networks*

Jonghyun Kim and Stephan Bohacek
Department of Electrical and Computer Engineering,
University of Delaware
Newark, DE, USA
jonghyun@udel.edu, bohacek@udel.edu

ABSTRACT

While opportunistic packet forwarding has been extensively studied in the case of single hop communication, only recently has opportunistic forwarding for multihop networks been explored. One key aspect of opportunistic forwarding is that the node that forwards a packet is determined on-the-fly and depends on which nodes receive the packet. On the other hand, in deterministic forwarding, the node that will forward the packet is predetermined. There are several key differences between these approaches. Opportunistic forwarding is able to make use of links where packet delivery is possible with a low probability and can easily accommodate channel fluctuations. But in order to reduce collisions, overhead is required to determine which node will forward the packet. The deterministic scheme considered is able to adaptively react to changes in the route quality, but requires overhead to update the route. In this paper, we examine which approach provides the best performance. It is found that the relative performance depends on mobility and the physical layer. However, in most cases, deterministic forwarding, along with route quality monitoring and high frequency route updates, outperforms opportunistic forwarding.

Categories and Subject Descriptors

C.2.1 [COMPUTER-COMMUNICATION NETWORKS]: Network Architecture and Design — *wireless communication*

General Terms

Performance

Keywords

Wireless, MANET routing, opportunistic forwarding

*This work was prepared through collaborative participation in the Collaborative Technology Alliance for Communications and Networks sponsored by the U.S. Army Research Laboratory under Cooperative Agreement DAAD19-01-2-0011. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiOpp'07, June 11, 2007, San Juan, Puerto Rico, USA.
Copyright 2007 ACM 978-1-59593-688-2/07/0006 ...\$5.00.

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, at the network layer there has been extensive effort focused on finding and maintaining routes (e.g. [8], [9], [13]). One approach is to exploit path diversity by switching between paths so that packets are forwarded along good paths.

There are several ways to exploit path diversity in multihop networks. One approach is to perform high frequency route updates [11], [14]. In this way, as the quality of routes degrade or as better routes become available, the route can be updated, and data packets will typically be transmitted over high quality routes. Another approach to exploiting diversity is to allow a set of nodes to act as relays. In this case, when a packet is transmitted, the best suited node will relay the packet. Such methods are known as opportunistic methods and are fundamentally different from the previously mentioned methods, which can be referred to as deterministic methods. The essential difference between the approaches is that opportunistic methods do not plan which node will forward the packet, while deterministic methods select the next hop before transmission. By postponing the decision as to which node will forward the packet, opportunistic methods are able to take advantage of random and, perhaps, rare opportunities. For example, due to the probabilistic nature of wireless transmission, a transmission may reach a node that is quite distance from the transmitter and close to the final destination. While this node might not be a reliable next hop, when the chance arises, it might be considerably better suited to forward the packet than a node that has reliable communication with the sender.

This paper compares the performance of opportunistic forwarding and deterministic forwarding with route quality monitoring and high frequency route updates¹. The basic question is whether the path diversity that arises due to the mobility and the related variability of radio links can be best accommodated through deterministic methods or through opportunistic methods. The paper especially focuses on the case where nodes are mobile. Furthermore, the impact of the radio is examined. Specifically, it is found that the performance depends on the “steepness” of the relationship between SNR and packet error probability. The

¹Deterministic forwarding with route quality monitoring and high frequency route updates will often be referred to as deterministic forwarding.

comparison is based on simulations. However, the simulations are realistic in the sense that the propagation is from ray-tracing computations and the mobility is from a realistic vehicle mobility model [7]. On the other hand, the impact of interference is not examined.

A brief summary of the results is that when nodes are mobile, deterministic forwarding along with route quality monitoring and high frequency route updates provides better performance than opportunistic forwarding. If the radio is such that the packet error probability decreases slowly as the SNR increases, then the difference in performance is smaller, but in terms of most metrics, deterministic forwarding remains the same or better than opportunistic forwarding. On the other hand, when nodes are not mobile and the packet error rate decreases slowly as the SNR increases, then opportunistic forwarding provides some benefits over deterministic forwarding.

2. OPPORTUNISTIC FORWARDING

A hallmark of opportunistic forwarding is that when a packet is transmitted, the next hop is not known, but depends on the outcome of the transmission. The motivation for this flexibility is that it allows multiple nodes to act as relays. For example, if the transmission fails to reach the primary relay, then a backup node may receive and relay the packet. In this way, opportunism increases reliability.

Opportunistic forwarding is also able to improve other performance metrics such as hop count and bit-rate. To see how this is possible, recall that in traditional, deterministic forwarding, a link is selected to forward data packets if it makes good progress toward the desired destination, and can support the desired bit-rate with an acceptable probability of success. Consequently, deterministic forwarding neglects links that can deliver packets closer to the desired destination, but that have large transmission error probability. However, if, by chance, the transmission across such a link succeeds, opportunistic forwarding is able to make use of the success. As a result, opportunistic forwarding can make use of links with high error probability, effectively decreasing the hop count and/or increasing the bit-rate.

The performance increase offered by opportunism is compounded when the node density is high, and hence, a large number of nodes are available to act as relays. Specifically, if the probability of successful transmission to node i is F_i , then the probability that one of the relays will receive the

packet is $1 - \prod_{i=1}^N (1 - F_i)$. As a result, even if F is small,

$1 - \prod_{i=1}^N (1 - F_i)$ may be quite close to one if N is large.

While opportunism offers great potential, there are several challenges, especially when nodes are mobile, which is the focus of this investigation. One important challenge is to determine exactly how to take advantage of the multiple relays. One possibility is to allow any node that receives the transmission to forward the packet. A problem with this approach is that these multiple transmissions may interfere with each other. An alternative approach is for the node that is best suited to deliver the packet to the destination to relay the packet. More specifically, the packet should be relayed by the node that is best suited *among all nodes that have successfully received the packet*.

The problem of which node is best suited to relay a packet toward the destination is the basic task of routing and can be addressed in several ways. For example, in [3], the suitability of a node to relay a packet was based on the geographic distance from the destination. The route metric used in this study is discussed in Section 3.

A more challenging problem is determining the set of nodes that have successfully received the packet. In [3], the following approach was used. The header of the data packet includes an ordered list of nodes that may relay the packet. The first node in the list has the highest priority; it is the best suited to relay the packet, but, perhaps, is the least likely to successfully receive the packet. After the data transmission completes, all nodes in the list listen to the channel to determine whether the node with highest priority is forwarding the packet. If, after a brief time period, no transmission is heard, the node that is second in the list has a chance to transmit. If this node fails to transmit, the next node on the list has a chance, and so on.

In [4], a similar method is used. However, in this case, no priority list is provided. Rather, when a packet is received, a node selects a random time that is based on its suitability to relay the packet, in particular, the more suitable, the shorter the mean time to transmit. Then, if no transmission is heard after waiting the random time, the packet is transmitted.

A drawback of these approaches is that each potential relay must be within communication range of every other potential relay, otherwise multiple nodes might transmit simultaneously, resulting in collisions. The requirement that potential relays must be within communication range reduces the number of potential relays, and hence reduces the degree to which diversity can be exploited. In the mobile environment, a more serious drawback is that even if potential relays can be selected so that they are all within communication range of each other, as nodes move, the links between relays may break, resulting in collisions. Hence, as the number of potential relays grows, the possibility of links breaking increases, and hence collisions become more prevalent. In early simulations, we found that the collisions were quite severe, and hence such approaches were not pursued.

We propose an alternative that does not suffer from collisions, but does require overhead. Like [3], the data packet includes an ordered list of relays. However, if the relay that is highest priority receives the packet, it transmits an ACK to the sender. This ACK is transmitted at a low data rate, hence, the sender will receive the ACK with a high probability. Upon receiving the ACK, the sender transmits an ACKACK. Again, this small packet is transmitted at a low data rate so that all potential relays receive the frame, and hence become aware that a relay with higher priority has received the data packet.

While this technique does require overhead, the overhead is one frame less than the RTS-CTS-DATA-ACK combination used in 802.11. On the other hand, if deterministic forwarding is used, then a data packet may be transmitted without a RTS or CTS. Furthermore, the transmission of the data packet by the downstream node can act as an implicit ACK. Hence, the proposed opportunistic scheme requires two more packet transmissions than the deterministic scheme that is used in this study. An important question answered in this investigation is whether the potential improvement in the performance offered by opportunistic forwarding is worth the extra overhead.

3. DETERMINISTIC FORWARDING WITH ROUTE MONITORING AND HIGH FREQUENCY ROUTE UPDATES

A distinguishing characteristic of deterministic forwarding is that there is a single route. As a result, the quality of this route can be easily monitored. Therefore, if the quality of the route degrades, then a route update can be performed. As a result, data packets will likely travel over high quality paths that provide high bit-rate and high packet delivery probability. One drawback of this approach is that overhead is required to update the route, and, of course, the forwarding is deterministic, and so the advantages of opportunistic forwarding are not available. In this section, techniques to monitor the quality of the route and update the route are discussed. Many aspects of the monitoring and route update are borrowed from a routing protocol that was designed to exploit diversity. More details on this protocol can be found in [11].

To judge the quality of the route, we employ a route metric referred to as the *worst-SNR-to-go to the destination* (or, for short, the worst-SNR-to-go). A route from a node to the destination is composed of a series of links. When a packet is transmitted over a link, the packet is received with a particular average SNR; we call this the SNR of the link. The worst-SNR-to-go to the destination from the node is the minimum of the each link's SNR along the path. In [6], it was found that the SNR is well correlated with link lifetime. Hence, the worst-SNR-to-go provides an indication of the route lifetime.

The worst-SNR-to-go from the source to the destination can be monitored as follows. When the source transmits a packet, it sets a worst-SNR-to-go field in the packet header to $+\infty$. When the next hop receives the packet, it measures the SNR during packet reception, and updates the worst-SNR-to-go with this recorded SNR and forwards the packet. When the next node receives the packet it compares the SNR during packet reception to the worst-SNR-to-go in the header and updates the packet header with the minimum of the two values. Hence, when the packet arrives at the destination, the worst-SNR-to-go from the source to destination is known.

One important use of path monitoring is to detect when the path quality has degraded, and hence a route update should be performed. To this end, the destination records the worst-SNR-to-go of the first packet received along a newly found or updated path. The destination then monitors the quality of the path upon every packet reception. If the quality of the path is found to degrade significantly below that of the first packet, then a route update is triggered.

3.0.1 J-Broadcast and efficient route updates

In [11], an efficient technique for performing localized route updates was developed. For completeness, the main aspects of this technique are repeated here.

When a route update is triggered, the destination broadcasts a J-broadcast (JBC) packet. All nodes that have overheard the data transmissions associated with this path will receive and process the JBC. Upon receiving the JBC, each node makes a determination as to whether the quality of the path to the destination is high enough. This determination is further discussed in Section 4. If the path is found to be suitable, then the node updates the worst-SNR-to-go field

in the JBC. The node then broadcasts the JBC. However, the JBC is broadcasted with CDMA and makes use of power control such that the transmission power is proportional to the worst-SNR-to-go. As shown in [11], CDMA, along with this power control, allows multiple JBCs to be received by each upstream node, and also ensures that JBCs advertising good routes are received with a high probability.

As the upstream nodes receive the JBCs, they measure the SNRs to the different downstream nodes that broadcasted the JBCs. Thus, associated with each received JBC is the received SNR (i.e., the SNR from the node that received the JBC to the transmitter of the JBC²) and the worst-SNR-to-go from the transmitter of JBC to the destination. We denote these values by SNR_i and $J2G_i$, respectively, where the subscript denotes the transmitter of the JBC. With these values, the node computes the worst-SNR-to-go to the destination via the transmitter of the JBC. We denote this value by $JVia_i$ and compute it from $JVia_i = \min(SNR_i, J2G_i)$. When the node decides to broadcast a JBC, the node computes its own worst-SNR-to-go via

$$\text{worst-SNR-to-go} = \max_i JVia_i. \quad (1)$$

This value is included in the JBC broadcasted by the node. In [11] and in the deterministic forwarding studied here, each node selects the next hop to be the node that achieves the maximum in (1).

This process of receiving JBCs, determining the worst-SNR-to-go, and broadcasting JBCs continues until the source receives JBC packets, at which point the source can determine a new path to the destination.

There are two key features of this approach to route update that results in good performance. First, the CDMA with power control greatly reduces the duration that the channel is busy as compared to a flooding-based route update. Also, while the near-far effect (see page 568 and Appendix E in [12]) may result in high loss probability when the received signal strengths greatly differ, the power control is able to take advantage of the near-far effect to filter out JBCs that advertise low quality routes. Second, since only nodes that have heard data packet transmissions will respond to the JBCs, the route update is localized.

4. PROTOCOL DETAILS

The purpose of this investigation is not to develop a routing protocol, but to compare opportunistic forwarding and deterministic forwarding. Thus, the protocols discussed here are focused on providing a fair comparison between opportunistic and deterministic forwarding, and should not be viewed as a routing protocol to be used in general settings.

For this comparison, we begin with a route found by AODV [8], i.e., RREQ flooding with RREP. Once a route is found, a single probe packet is sent along the route. All nodes that are within communication range of any node along the path become aware of the path and prepare to receive and process JBC packets. When the probe packet reaches the destination, the destination broadcasts a JBC

²According to reciprocity (e.g., page 90 in [2]), the SNR over a link is independent of which link is transmitting and which link is receiving. On the other hand, SNIR is dependent on which end of the link is transmitting. However, the impact of interference on opportunistic and deterministic forwarding is outside of the scope of this investigation.

packet. The propagation of the JBC packet is similar to what was described in Section 3. However, Section 3 does not mention how a node decides whether to forward a JBC; this decision depends on whether deterministic or opportunistic forwarding is deployed and is discussed next.

In both approaches, as JBCs are received, the node computes the probability of successful transmission to the downstream nodes. In the case of deterministic forwarding, this probability is

$$\max_{i \in \text{received JBCs}} F_i(BR_0), \quad (2)$$

where $F_i(BR_0)$ is the probability of successful transmission to downstream node i when the packet is transmitted at the lowest bit-rate, BR_0 , and the maximization is overall downstream nodes from which a JBC has been received. In the case of opportunistic forwarding, the probability of successful transmission is

$$1 - \prod_{i \in \text{received JBCs}} (1 - F_i(BR_0)). \quad (3)$$

In either case, if the probability of successful transmission exceeds Target Transmission Prob. (0.999 was used in this investigation), then the node prepares to broadcast a JBC. Otherwise, the node continues to collect JBCs. Note that for a given set of received JBCs, the opportunistic forwarding will have a larger probability of successful packet transmission, and hence, opportunistic forwarding might allow the transmission of JBCs sooner than deterministic forwarding. This would result in a smaller hop count for opportunistic forwarding.

Once the probability of successful transmission is above the threshold, the node waits for the current set³ of JBCs to be received. The node sets its hop count to be one plus the maximum hop count of all collected JBCs. In the case of deterministic forwarding, this hop count matches the number of hops that a packet will traverse to the destination. In the case of opportunistic forwarding, a packet might traverse fewer hops than the hop count. Nonetheless, this hop count metric is a useful metric to gauge how suitable the node is for delivering a packet to the destination.

Next, the node computes $JVia_i$ and computes its own worst-SNR-to-go, as described in Section 3. In addition, and regardless of whether deterministic or opportunistic forwarding is used, the node that achieves the maximization in (1) is determined. We refer to this node as the *target node*, i.e.,

$$\text{target node} = \arg \max_i JVia_i.$$

In the case of deterministic forwarding, the target node is set as the next hop. However, in the case of opportunistic forwarding, the role of the target node is more complex as discussed next.

The goal of opportunistic forwarding is to take advantage of the opportunities that a packet is received by a node that is closer to the destination than is expected by the deterministic forwarding. That is, opportunistic forwarding should take advantage of receptions by nodes that are better suited to deliver the packet to the destination than the target node. We classify nodes as follows.

³JBCs tend to be transmitted in bursts, where a set of nodes that are the same number of hops from the destination transmit the JBCs together.

DEFINITION 1. A preferred node is a node that has a smaller hop count than the target node or has the same hop count but a larger $J2G$.

DEFINITION 2. A backup node is a node that has the same hop count as the target node but a smaller $J2G$.

Note that a node cannot have both a larger $J2G_i$ and SNR_i than the target node. Consequently, the preferred nodes have smaller SNR_i than the target node. Also, the backup nodes typically have a larger SNR_i than the target node. Hence, the data packets are less likely to reach preferred nodes than they are to reach the target node, and data packets are more likely to reach the backup nodes than to reach the target node.

In opportunistic forwarding, the packet header includes an ordered list of nodes with the first node having the highest priority for forwarding the packet. The ordering is as follows

1. A node with lower hop count has a higher priority than a node with higher hop count.
2. If two nodes have the same hop count, then the node with the larger $J2G$ has higher priority.
3. If two nodes have the same hop count and the same value of $J2G$, then the node with larger $JVia$ has higher priority.
4. Further ties are broken randomly.

The target node, the preferred nodes, and the backup nodes are all listed as potential relays in the data packet header.

The final task that a relay must perform before it is ready to forward packets is to determine the bit-rate to transmit data packets. In the case of deterministic forwarding, the bit-rate is determined from

$$\begin{aligned} \text{bit-rate} &= \max BR \\ &\text{such that: } F_{\text{TargetNode}}(BR) \\ &\geq \text{Target Transmission Prob.} \end{aligned} \quad (4)$$

where $F_{\text{TargetNode}}(BR)$ is the probability of successfully transmitting a packet to the target node while transmitting at bit-rate BR . Note that $F_{\text{TargetNode}}(BR)$ depends on the SNR to the target node, which is known.

In the case of opportunistic forwarding, the bit-rate is computed as follows

$$\begin{aligned} \text{bit-rate} &= \max BR \\ &\text{such that: } 1 - \prod_{\substack{i \in \text{preferred node(s)} \\ \text{or target node}}} (1 - F_i(BR)) \\ &\geq \text{Target Transmission Prob.} \end{aligned} \quad (5)$$

Hence, since opportunistic forwarding takes advantage of all potential relays, it may be able to transmit at a higher bit-rate than the deterministic case. Note that the bit-rate is determined by only considering the target node and preferred nodes. The backup nodes are not considered. The reason for this is that the backup nodes may be very close to the sender. Hence, these links may support very high bit-rates. However, as compared to the target node, these relays do not make good progress toward the final destination. If these nodes were included into the bit-rate calculation, the

maximum bit-rate would be set to reach these backup nodes and hence the preferred nodes would have little chance to receive and forward the data packet.

Once the above calculations are performed, the node broadcasts the JBC packet with CDMA and power control as described in [11]. The source is ready to forward packets once it has received enough JBCs such that (2) exceeds Target Transmission Prob. or (3) exceeds Target Transmission Prob, depending on if deterministic or opportunistic forwarding is being used.

During transmission of the data packets, both deterministic and opportunistic forwarding make use of channel measurements. In the case of deterministic forwarding, the route is monitored and if it is found to degrade below a threshold, then a route update is triggered. Furthermore, when a node transmits a data packet, the upstream node is able to measure the SNR to the transmitting node. This permits a recomputation of the bit-rate via (4) or (5).

4.1 Mixed forwarding - opportunistic forwarding with route monitoring and updating

In order to gain further insight into the differences between deterministic and opportunistic forwarding, a mixture of the two is considered. We refer to this technique as mixed forwarding. Mixed forwarding attempts to use the best features of opportunistic forwarding and the best features of deterministic forwarding. Mixed forwarding uses opportunistic forwarding in that each data packet has an ordered list of potential relays. The ordering of the list and the transmission bit-rate is selected in the same way as opportunistic forwarding explained above. Furthermore, in order to determine which node will relay the packet, the Data-ACK-ACKACK scheme discussed in Section 2 is used. However, instead of solely relying on opportunism, the destination monitors performance of the forwarding. Specifically, if the destination determines that a backup node has forwarded the packet, then a route update is initiated.

While the mixed forwarding case allows packet reception opportunities to be exploited and it allows the route to be monitored and updated, it also suffers from the overhead of the ACK-ACKACK and from the JBCs to update the route. Hence, there is a possibility that the mixed case may perform worst than both deterministic and opportunistic forwarding.

5. SIMULATED ENVIRONMENT

5.1 Mobility and propagation

The mobility and propagation environment play an important role in the performance of routing protocols. The simulated environment here is based on realistic urban propagation and mobility. The propagation is based on ray-tracing computation of a 2 km \times 2 km region of downtown Chicago. As a result, the propagation is realistic. The mobility is based on a realistic vehicle mobility model. The model accounts for traffic lights and internode dynamics such as passing. More details on the propagation and mobility can be found in [7]. The simulations model midday mobility, i.e., not rush hour or lunchtime. Five node densities were examined, specifically, 64, 128, 256, 512, and 1024 nodes. These scenarios are number 1 through 5, respectively. In all cases, a 512 byte data packet was sent once every 50 msec.




Figure 1: Left: The PEP/SNR relationship for all six types of radios considered. Right: The relationship between PEP and SNR for three steepness and two bit-rates.

Each connection lasted 5 minutes. 60 source-destination pairs were simulated. The sources and destinations were the same regardless of the node density. Packet simulations were performed with QualNet.

5.2 Radio model

Given the location of two nodes, the propagation model determines the channel between the two nodes. The probability of correctly decoding the packet depends on the radio, or more specifically, on the modulation and the code. The radio plays an important role in the performance of networks in general, and plays a similarly important role in the relative performance of deterministic and opportunistic routing. This investigation examines the performance under a sequence of radio models. All models are based on M-ary modulation without error correction coding. Hence, the nominal radio model follows the standard model of M-ary transmission error probability (e.g., page 327 in [12]). Bits per symbols considered include 1, 2, 4, 6, 8, 10, and 12.

In order to model a variety of radios, we adjust the steepness of the relationship between the packet error probability (PEP) and SNR. Figure 1 shows some relationships between PEP and SNR. In this investigation, we consider steepnesses that range between steepest, steep, nominal, shallow, shallower, and shallowest (for lack of better terms). The nominal relationship is directly from [12]. In the steepest case, the packet is received either with probability one or zero. This is similar to an infinite length code such that if the SNR exceeds a particular threshold, then perfect data reception is achieved, and no data is correctly decoded if the SNR is below the threshold. In the shallowest case, there is only a moderate relationship between PEP and SNR.

Note that in Figure 1, the SNR where the PEP is 10^{-3} remains the same regardless of the steepness of the PEP/SNR relationship. This idealization of the radio allows the impact of the steepness of the PEP/SNR relationship to be investigated in isolation from the other impacts that changing the modulation and code may have. For example, keeping the point where PEP is 10^{-3} fixed results in the connectivity not being greatly impacted by the steepness of the PEP/SNR relationship. Indeed, if there exists a path between a source and destination such that each link has a PEP of less than 10^{-3} with one radio, then this same route is available with the other radios. Consequently, this investigation can focus on forwarding, and not on route finding. However, it should be emphasized that we make no claim that the behavior displayed in Figure 1 is typical for changing modulation and codes.




Figure 2: The bit-rate of the first packet sent. Note that the legend is used for the rest of plots.

An important aspect of the radios is that in the steepest case, if the channels are known (as they are in the case of deterministic forwarding), then no opportunism is possible. That is, if the channels are known, then the transmitter knows exactly which nodes will receive the packet. To put it another way, opportunistic forwarding takes advantage of events that cannot be predicted with certainty. However, if the PEP/SNR relationship is steep, then all outcomes are predictable. On the other hand, if the PEP/SNR relationship is shallow, then it is not possible to predict which nodes will correctly receive transmissions, and hence, there are events that opportunism may exploit.

6. SIMULATION RESULTS

6.1 The first packet (before nodes move)

We begin with an examination of the bit-rate shown in Figure 2. This figure as well as all figures shows the results from 60 trials in each scenario. There are 30 scenarios, corresponding to 5 node densities (from 64 nodes to 1024 nodes), and 6 PEP/SNR relationships.

Figure 2 shows the average bit-rate of the first packet. As discussed in Section 4, the route is found by flooding, sending a probe packet, and then sending JBC packets that searches for a good route. Deterministic and opportunistic forwarding perform this search in a similar way, and hence the routes used are nearly the same. In particular, when the PEP/SNR relationship is very steep, deterministic and opportunistic forwarding use the exact same route and the same bit-rate. However, as the PEP/SNR relationship becomes shallower, opportunistic forwarding permits a higher bit-rate. This improvement in bit-rate is extended as the node density increases. Since mixed forwarding uses the same method to select the bit-rate as opportunistic forwarding, it provides the same bit-rate. Hence, Figure 2 clearly demonstrates opportunism at work, and also shows how opportunism is impacted by node density and the PEP/SNR relationship. It is interesting to note that if the PEP/SNR relationship is nominal or steeper, then opportunism provides a bit-rate that is approximately the same as deterministic forwarding. However, if the PEP/SNR relationship is shallower than nominal, opportunism provides a bit-rate that may exceed the bit-rate provided by deterministic forwarding by approximately 50%.

Note that Figure 2 only examines the first data packet. Hence, the impact of mobility is not shown. This performance is similar to what might be observed in a stationary wireless network such as a mesh network.

Figure 3 provides another glimpse of opportunism. This figure shows the received signal strength of the data packet by the node that relays the packet. Clearly, opportunistic forwarding shows a weaker signal strength. Indeed, one of the key features of opportunistic forwarding is the ability to reliably transmit over links that have a low channel gain. Figure 3 also illustrates an important difference between opportunistic and deterministic forwarding. In the case of deterministic forwarding, when the received signal strength decreases, it is an indication of a decrease in the route quality. Hence, the received signal strength can be used to determine route quality. However, in the opportunistic case, such measurements are not possible. In fact, it is hoped that the received signal strength is low. As a result, when opportunistic forwarding is used, it is difficult to also monitor the route quality and perform route updates.

Figure 3 clearly shows how as the PEP/SNR relationship becomes more shallow, opportunism can be exploited. This differs from Figure 2 where the three steepest radios performed the same. The reason for this is that only a discrete set of bit-rates were examined. For example, if, instead of bit-rates 1, 2, 4, ..., we considered bit-rates 1, 1.25, 1.5, ..., then all radios would display distinct bit-rates. That is, the equality of the bit-rates for the nominal and steeper radios is due to quantization of the bit-rate. Another impact of such quantization is shown in Figure 4, which shows the average number of hops. This figure appears to be a single line, however, it turns out that all techniques and all radios provide the exact same number of hops.

6.2 Performance before the first route failure

Now we focus on the data packets sent after the first packet, but before any of the schemes experiences a route failure. Note that since the schemes will experience route failures at different times, focusing on the time before any scheme has a route failure allows the paths to be approximately the same. This also allows the comparison to be focused on forwarding, not on routing issues such as route discovery and rediscovery.

Figure 5 shows the average bit-rate until the first route failure. This figure shows that the improvement displayed by opportunism in Figure 2 disappears when mobility is considered. Furthermore, the bit-rate offered by deterministic forwarding surpasses that of opportunistic by no more than 35% and is typically between 5% and 10% better. The mixed forwarding provides a bit-rate that is between opportunistic and deterministic. To understand the reason for the behavior, recall that in all approaches when a downstream node transmits a packet, the upstream nodes record the SNR and based on this measurement, the bit-rate is updated. Therefore, in the case of opportunistic forwarding, if nodes move, the SNR may decrease and the bit-rate is decreased. On the other hand, when the SNR decreases in the deterministic case, the destination detects the degradation in the route and triggers route update that discovers a better route. The deterministic forwarding performs slightly better than the mixed case because, as mentioned above, deterministic forwarding is better able to monitor the path quality.

Figure 6 shows the average number of hops traversed by data packets before the first route failure. While Figure 4 shows that the initial routes were the same length, as nodes move, this behavior changes. Deterministic and opportunistic forwarding change path lengths in different ways. In the




Figure 3: The average signal strength received by the relay that transmits the packet.




Figure 4: The average number of hops that packets traverse. This plot is for all scenarios and approaches.




Figure 5: The average bit-rate before the first route failure.

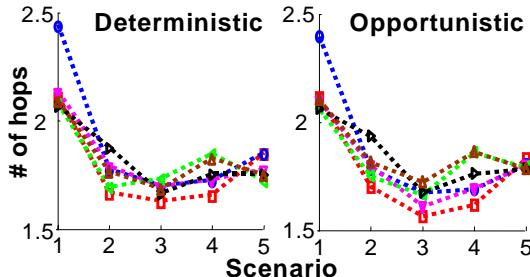


Figure 6: The average number of hops before the first route failure.

opportunistic case, a packet may traverse fewer hops if a preferred node is able to decode the packet. Indeed, one important benefit of opportunistic forwarding is the ability to take advantage of such events. On the other hand, deterministic forwarding makes use of shorter routes because it performs route updates and explicitly finds a shorter route. While these two approaches are quite distinct, remarkably, the resulting average hop counts are quite close; on average, deterministic forwarding uses a route that is 0.5 % shorter. Note that Figure 6 does not show the mixed forwarding case; it also has the same average hop count.

6.3 Performance during the connection lifetime

A common routing performance metric is packet delivery ratio. Figure 7 shows the packet delivery ratio for the three approaches. The data used to make this figure is from the entire five-minute connection. This figure shows that the repeated route updates performed by deterministic forwarding is more critical for maintaining a path than allowing opportunistic forwarding. While the difference in performance is decreased as the PEP/SNR relationship becomes less steep, deterministic forwarding remains better than opportunistic forwarding. The mixed case shows performance that is similar to the deterministic forwarding, providing further indication that the repeated route updates are critical for packet delivery.

While packet delivery ratio is a useful metric, in these experiments, most packet delivery failures were due to route failures. Hence, to better understand packet delivery, we consider the route failure rate shown in Figure 8. This figure mimics Figure 7. There are two reasons that opportunistic routing experiences shorter route lifetimes than the deterministic forwarding examined here. First, determinis-




Figure 7: The packet delivery ratio.

tic forwarding is able to perform route updates when the path degrades. Since it is common that the route quality degrades before it fails [6], path monitoring is critical to maintaining connectivity.

Second, opportunistic forwarding is actually more sensitive to link failure than deterministic forwarding. To see this, consider (5). Opportunistic forwarding is able to support a high bit-rate since there are many links that it may utilize. While such an approach is sensible when nodes are not moving, it leads to short route lifetime when nodes are mobile. Specifically, the bit-rate provided by (5) assumes that all links are in place. Thus, while the ability of deterministic forwarding to forward a packet depends on a single link, the ability of opportunistic forwarding to forward a packet depends on many links. If *any* of the links decreases in quality, then the effective transmission probability error will rise above its target value, and may result in a route error. In short, in order to improve performance, opportunistic forwarding becomes dependent on more links, and hence is more likely to suffer from a route failure.

A by-product of this sensitivity is that the backup nodes are forced to relay packets. In the case of mixed forwarding, when a backup node forwards the packet, a route update is triggered. Hence, part of the reason that the mixed for-




Figure 8: The average rate of route failures.

warding approach provides good performance is due to its frequent route updates, which is due to the sensitivity to mobility. Furthermore, our initial version of the opportunistic test protocol did not make use of SNR measurements from downstream nodes to readjust the bit-rate. However, this lack of reactivity exasperated this sensitivity to mobility, and hence, in order to make a fair comparison, opportunistic forwarding was made more reactive.

Finally, Figure 9 shows the overhead efficiency. We define overhead efficiency as the duration that data packets are transmitted divided by the duration that any packet including overhead is transmitted. Note that duration is used, hence we measure how efficiently the channel is being utilized. Figure 9 shows that opportunistic forwarding suffers from considerably lower efficiency as compared to deterministic forwarding. Thus, while deterministic forwarding makes use of route updates, these updates are not particularly frequent and are highly efficient (recall that multiple JBCs are sent simultaneously, efficiently using the channel). On the other hand, opportunistic forwarding requires the Data-ACK-ACKACK sequence of frames; the extra ACKs are sent for each packet and reduce the efficiency.

7. CONCLUSIONS

Two distinct network layer techniques were examined, opportunistic forwarding and deterministic forwarding with route monitoring and route updates. It was found that when there is no mobility and when the PEP/SNR relationship is shallow, opportunistic forwarding provides some performance improvements over deterministic forwarding. However, deterministic forwarding has lower overhead. When nodes are mobile, the performance benefits provided by opportunism evaporates. To some degree, a shallower PEP/SNR relationship improves the performance of opportunistic forwarding. This is expected since a shallow PEP/SNR relationship results in events that opportunistic forwarding can take advantage of. However, while a shallow PEP/SNR relationship does decrease the gap between deterministic and opportunistic forwarding, deterministic forwarding still outperforms opportunistic forwarding. Finally, to some extent, mixed forwarding provides many of the benefits of the other schemes. However, it also suffers from the high overhead as pure opportunistic forwarding does. Furthermore, it does not provide a substantial improvement over deterministic forwarding unless the PEP/SNR relationship is shallow and the nodes are not moving.

These results imply that carefully monitored and efficient route updates are more useful than opportunistic forwarding. To put it another way, it is well known that mobile multihop wireless networks provide path diversity [5], i.e., some paths are better than others. Furthermore, the path




Figure 9: The overhead efficiency.

quality is time varying. This investigation indicates that the type of diversity that occurs in *mobile* multihop wireless networks is better exploited through route monitoring and route updates than by opportunistic forwarding techniques.

It is important to note that this study did not investigate the impact of interference.

Disclaimer

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U. S. Government.

8. REFERENCES

- [1] J. R. Barry, E. Lee, and D. G. Messerschmitt. *Digital Communication*. Kluwer Academic Publishers, Boston, 2004.
- [2] H. L. Bertoni. *Radio Propagation for Modern Wireless Systems*. Prentice Hall, 1999.
- [3] S. Biswas and R. Morris. ExOR: Opportunistic multi-hop routing for wireless networks. In SIGCOMM, 2005.
- [4] A. Bletsas, A. Khisti, D. P. Reed, and A. Lippman. A simple cooperative diversity method based on network path selection. *IEEE JSAC* 24, pages 659–672, 2006.
- [5] S. Bohacek. Achievable performance improvements provided by cooperative diversity. *WiOpt*, 2006.
- [6] S. Bohacek, A. Ilic, and V. Sridhara. On the predictability of link lifetimes in urban MANETs. In WiOpt, 2005.
- [7] S. Bohacek, V. Sridhara, and J. Kim. UDel models.
- [8] E. M. R. Charles E. Perkins. Ad hoc on-demand distance vector routing. *2nd IEEE Workshop on Mobile Computing Systems and Applications*, pages 90–100, 1999.
- [9] D. B. Johnson, D. A. Maltz, and J. Broch. DSR: The dynamic source routing protocol for multi-hop wireless ad hoc networks. In AdHoc Networking, pages 139–172. Addison-Wesley, 2001.
- [10] J. Kim and S. Bohacek. Selection metrics for multihop cooperative relaying. In MedHocNet, 2005.
- [11] J. Kim and S. Bohacek. Exploiting multihop diversity through efficient localized searching with CDMA and route metric-based power control. In MSWiM, 2006.
- [12] T. Rappaport. *Wireless Communication*. Pearson Education, 2002.
- [13] C. Sanivanez and R. Ramanathan. Hazy sighted link state routing protocol (HSLs). Technical Report 1301, BBN, 2001.
- [14] P. Sambasivam, A. Murthy, and E. M. Belding-Royer. Dynamically adaptive multipath routing based on AODV. In MedHocNet, 2004.