Design Guidelines for Routing Protocols in Ad Hoc and Sensor Networks with a Realistic Physical Layer

Ivan Stojmenovic, Amiya Nayak, and Johnson Kuruvila, SITE, University of Ottawa

ABSTRACT

We present guidelines on how to design network layer protocols when the unit disk graph (UDG) model is replaced by a more realistic physical layer model. Instead of merely using the transmission radius in the UDG model, physical, MAC, and network layers share the information about a bit and/or packet reception probability as a function of distance between nodes. We assume that all nodes use the same transmission power for sending messages, and that a packet is received when all its bits are correctly received. The MAC layer reacts to this probabilistic reception information by adjusting the number of acknowledgments and/or retransmissions. We observe that an optimal route discovery protocol cannot be based on a single retransmission by each node, because such a search may fail to reach the destination or find the optimal path. Next, we discuss that gaining neighbor knowledge information with "hello" packets is not a trivial protocol. We describe localized position-based routing protocols that aim to minimize the expected hop count (in case of hop-by-hop acknowledgments and fixed bit rate) or maximize the probability of delivery (when acknowledgments are not sent). We propose a guideline for the design of greedy position-based routing protocols with known destination locations. The node currently holding the message will forward it to a neighbor (closer to the destination than itself) that minimizes the ratio of cost over progress, where the cost measure depends on the assumptions and metrics used, while the progress measures the difference in distances to the destination. We consider two basic medium access layer approaches, with fixed and variable packet lengths. This article will serve as a preliminary contribution toward the development of network layer protocols that will match the assumptions and criteria already used in simulators and ultimately in real equipment.

WHAT IS THE TRANSMISSION RADIUS?

The designers of network layer protocols for ad hoc and sensor networks assume the unit disk graph (UDG) communication model, where two nodes communicate if and only if they are within distance R, where R is the transmission radius, equal for all nodes. Almost all articles even use R as the independent variable in their simulations. While the protocols at the network layer are designed with simple assumptions and performance metrics, experiments are normally carried on simulators that implement more realistic physical and medium access control (MAC) layers.

Simulators are trying to match the physical layer, which suggests that the UDG model is not realistic because it ignores random variations in received signal strengths. It was demonstrated that signal strength fluctuations have a significant impact on ad hoc network performance metrics, sometimes "outperforming" the impact of node mobility. Thus, nondeterministic radio fluctuations cannot be ignored when designing robust ad hoc network protocols based on ad hoc network simulation and analysis.

Assuming fixed signal-to-noise ratio (SNR), the model used in simulators and that hopefully matters in real equipment then looks like the one in Fig. 1, which shows how packet reception probability $p(x)$ depends on distance x between two nodes. The exact shape of the curve depends on the exact model used (combined Friis and two-ray ground model in [1]; lognormal shadowing model [2, 3]). It is obvious that the UDG model is indeed a good initial approximation for this, since the reception probability is close to 0 or 1 everywhere except around the transmission radius. But what is the transmission radius in Fig. 1? Is $R = 30$ (refer to distances in Fig. 1), meaning that the failure rate for transmissions is ≤ 5 percent? Is $R = 50$, meaning that the reception probability is ≥ 5 percent (i.e., two nodes

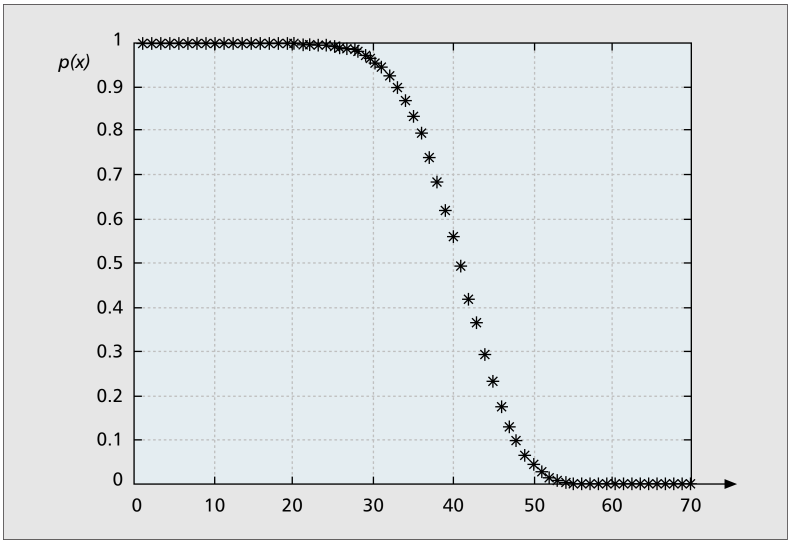

Figure 1. *Packet reception probability as a function of distance in a typical physical layer model.*

can still communicate if they make a reasonable number of attempts)? Is $R = 41$, meaning that the distance where the (bit or packet?) reception probability is 0.5?

Consider the impact of this decision on the greedy position-based routing protocol by Finn [4]. In this protocol the node currently holding the packet will forward it to the neighbor closest to the destination (routing fails if no neighbor is closer to the destination than the current node). The metric used to measure optimality is hop count: the number of links on the created route. Neighbors are nodes that are defined by the UDG model. If a conservative decision about R is made ($R = 30$), one transmission normally suffices for each hop. The hop count metric is then reasonably accurate, but the network may become unnecessarily partitioned. Why is it that two nodes that can communicate with approximately two or three attempts each time are then ignored? If $R = 50$, the selected neighbor is likely to be the one that requires several retransmissions to forward the message. In that case, the hop count metric does not properly reflect the cost involved. The expected hop count (EHC) needs to be used instead, which measures all the transmissions on each hop. However, when this metric is used, the greedy routing protocol does not provide the optimal solution [5]. The reason is that it is a threshold-based approach, favoring distant neighbors that require a significant number of retransmissions to be reached, thus causing large EHC.

The exact functions $p(x)$ in specific models used in literature are quite complex and time consuming to compute. We proposed [5] to use instead a reasonably accurate but simple *approximation* of the function. When the packet reception probability is used, it is calculated based on the bit error rate (complementary to bit reception probability) and packet length. The reason to use the approximation rather than the actual function is to reduce computation time at each node in both simulations and in real equipment, and to simplify the analyses and simulation of it.

We also defined two nodes as being neighbors if the packet reception probability between them is above a certain minimum threshold (we used a threshold of 0.05). For simplicity and convenience of describing protocols, the transmission radius R is assumed to be the distance where $p(R) = 0.5$. For example, the neighbors are then nodes at distances $\leq 1.44R$ for power attenuation $\beta = 2$ and packets with 120 bits.

It is important to observe that our proposed protocols remain the same under any realistic physical layer model. They only require a bit/packet reception probability function $p(x)$ to be a shared parameter across layers. Nodes can estimate the probability of receiving a bit or packet based on either signal strength, distance between nodes, or merely by deriving statistics from a number of bits or packets sent and received recently between two nodes. Lower layers therefore decide what the function $p(x)$ is and pass it to the upper layers, which use it to decide, for example, what the best forwarding neighbor is in position-based routing. The number of attempts to reach a neighbor at the medium access layer also depends on $p(x)$, which then also defines the EHC on a given link.

The decision on whether to use bit or packet reception probabilities depends on the assumptions made. In our preliminary work [5, 6] we divided the message into fixed size packets, with acknowledgments of the same length; therefore, packet reception probability was an appropriate choice. If one assumes variable packet sizes at the MAC layer (as used in [1]), the bit reception probability is more appropriate for use (however, the transmission radius, for the purpose of defining neighbors, needs to be redefined again to take into account maximal or average packet lengths).

This article addresses the implications of replacing the UDG model with the model where bit/packet reception probability depends on distance between nodes for the design of basic network layer protocols: broadcasting or route discovery, "hello" packets, and position-based routing. Technical details of the proposed solutions (will) appear elsewhere (e.g., [3, 5, 6]). The main contribution here is to describe some design principles and new paradigms involved, thereby moving closer to reality in the basic model. We anticipate that our design principles will help in addressing and solving other milestones toward designing efficient network layer protocols, verified by real equipment.

Route Discovery in Ad Hoc Networks

In reactive routing protocols, nodes follow the flooding based route discovery procedure (e.g., DSR or AODV [4]) to find the route to the destination. During the route discovery process, each node receiving the route message will find the cumulative cost of the route from the source to itself (measured as hop count, power consumption, or any other metric), and decide whether the cost is lower than the best cost it has received so far for the same task. If so, it will retransmit the message *exactly once*. The destination will report back to the source using the memorized best path. A distant neighbor may accidentally receive the only route discovery message and happily use it to announce a new route, although in reality it is not able to provide such a good service for the real traffic. Another issue is that only one communication can miss an important neighbor and therefore the route may not be found. Overall, the constructed route may be far from the optimal one. Morever, the route constructed by control traffic may not be available to data traffic because of difference in packet lengths.

Several studies already observed that reactive routing protocols do not perform well when a realistic physical layer is considered. An MIT group [7] proposed to use the expected transmission count metric (ETX) to send high-throughput paths on multihop wireless networks. The ETX metric takes into account the effects of link loss ratios, asymmetry in the loss ratios between the two directions of each link, and interference among links of a path. Then they apply the ETX metric to reactive routing protocols and show that the ETX metric improves performance. Their observations are based on a real implementation, without giving any theoretical results or analysis in support.

Qin and Kunz [2] concentrate on the impact of a realistic physical layer (shadowing propagation model) on simulating the performance of the well known AODV and DSR on-demand wireless routing protocols. They proposed new signal power thresholds for route discovery to enable the selection of links with strong enough signal strength and reduce some protocol control messages. They report a significant increase in the packet delivery ratio and a decrease in packet latency, and suggest that link status is a better metric than hop count for selecting routes in shadowing models. We believe that the *best route* approach should be used instead of the *threshold*-based one, since thresholds may prevent operational links from being included in possibly the best or sometimes the only route.

Nadeem and Agrawala [1] define energy efficiency as the expected number of bits for sending a packet between two nodes, which is then extended to cumulatively measure energy efficiency of a route. In their approach a message is divided into packets of variable size. The fragmentation of messages into variable size packets is supported by the well adopted IEEE 802.11 MAC layer protocol. The optimal packet length is described by a formula that depends on the bit error rate, number of bits in the message overhead, and number of preamble bits transmitted with each packet [1]. After all packets are accumulated at the receiver, a new optimal packet size for the next hop is determined. They assume that acknowledgments are packets with a much smaller number of bits and do not consider them in the cost to simplify the derivation.

The common problem in all of the described attempts [1, 2, 7] is that each node still retransmits the received packet (at most) once. While blind flooding is an acceptable protocol in the UDG model, guaranteeing connectivity and optimality, it may not lead to optimal or even any solution under the given realistic physical layer. A single transmission may not reach a particular neighbor, which then can disconnect the destination from the source or fail to offer a good existing route. A link that failed in only one attempt between two nodes (during a route discovery process) may be part of a good link and contribute toward an optimal route under the considered metric.

We therefore propose a modification to the route discovery process. A single packet retransmission by a node considered for inclusion may not be received by a relatively distant neighbor that is a good choice for a final constructed route. Therefore, each node may retransmit the given route discovery packet *several times* rather than once as in the UDG model. Such multiple retransmissions may also serve to measure or reevaluate the packet reception probability. What is the optimal number of retransmissions? A simple option is to retransmit a fixed number of times (e.g., twice or three times). In dense neighborhoods, a few retransmissions may suffice. Better trade-offs between messages sent and gains made could be achieved when a node retransmits until a certain number of packets with the same content have been received, before or after the first retransmission (subject to a timeout) [3]. Nodes may also overhear further packets for the same route requests, and make decisions to stop based on satisfactory progress of the route request. If a received message contains information about a better route, the retransmission of the previously known best route (if any is ongoing) stops, and the retransmission of the new route starts with a fresh counter. One particular issue is the existence of long "bridge" edges between two subnetworks. The link may not be discovered unless a maximal number of attempts are made. To increase the chance of discovering such bridges, nodes that are not able to extend a currently offered partial path would continue retransmitting until the maximum allowable number of packets is sent. These general descriptions may be completed toward particular protocols in a variety of ways. The optimal number of retransmissions and therefore the cost of applying this concept primarily depend on the network density. In very dense networks, a single retransmission of route request may find a sufficient number of neighbors to provide close to the best route. However, the sparser the network, the more retransmissions might be needed, especially for discovering important bridges in the network. Network traffic conditions also impact the cost of finding near optimal routes.

The metric used during the route discovery process depends on the assumptions. For instance, if a route with hop by hop Acknowledgments is searched for, the appropriate metric is the expected hop count (EHC) on the route, as the sum of expected hop counts on each hop. The calculation of EHC further depends on medium access assumptions. In [5], a message is divided into packets of fixed size, and each packet is then independently routed. Acknowledg-

The optimal packet length is described by a formula that depends on the bit error rate, the number of bits in the message overhead, and number of preamble bits transmitted with each packet. After all packets are accumulated at the receiver, a new optimal packet size for the next hop is determined.

Note that in different scenarios, such as reasonably stable networks with periodic hello messages, single hello messages may perform well, since protocols do not need to immediately declare a link as down if a hello message is lost once in a while.

ments, if sent, are of the same size as the packet, and are counted in the expected hop count. EHC then depends on the expected number of retransmissions and expected number of acknowledgments.

The cost of a link on a route is estimated as follows. In [3, 5] it is observed that the optimal number of acknowledgments is $u \approx 1/p(x)$, which gives the best EHC $2/(p(x)(1 - (1 - p(x))^u))$. This cost can be included in the transmitted packets for addition to the considered route cost. Note that [5] actually derives the ideal hop count (if additional nodes can be placed in ideal positions) and uses it in timeout decisions for retransmissions. This allows better routes to be advertised faster than others, and consequently reduces overall traffic, since nodes that overhear neighbors already announcing themselves as better "forwarders" may cancel their own retransmissions.

EHC assumes a fixed bit rate. If the bit rate depends on distance, EHC needs to be replaced with the expected time to forward the packet. If medium access allows variable packet lengths and acknowledgments with few bits (as considered in [1]), EHC can measure the expected number of bits transmitted. In case of a variable bit rate and variable packet length, the expected time for transmitting all bits can be used.

In case of routing without acknowledgment, the optimality measure is the probability of delivery along a given route. This scenario may occur in sensor networks, where a sink establishes reverse paths to sensors, which may periodically report their data to the sink without relying on any acknowledgment. The probability of delivery is equal to the product of delivery probabilities along each link on a route.

We apply a similar idea of several retransmissions for gaining neighbor knowledge in ad hoc networks. We are primarily considering a mobile ad hoc network where nodes enter a new neighborhood, or the initialization phase in sensor networks. Most existing network layer protocols assume that each node is aware of its one-hop (often even two-hop) neighbors. To find out about the existence of neighbors, nodes need to resort to hello message exchange. This is a simple procedure in the UDG model, accomplished by each node sending one hello packet, which is then received by all neighbors located within transmission radius R. However, with a realistic physical layer, a hello message operation requires a closer look. A simple solution is that each node retransmits a hello packet a certain fixed number of times. To reduce the number of transmissions without significantly sacrificing connectivity, we proposed [3] a variable number of hello packets sent in order to learn about neighbors. In this protocol, packets are sent until a satisfactory number of neighbors is discovered (or timeout expires). In order to learn about all neighbors, which may be important for preserving connectivity, nodes may be forced to resend hello messages a maximal number of times. Note that in different scenarios, such as reasonably stable networks with periodic hello messages (e.g., in OLSR protocol), single hello messages may perform well, since protocols do not need to immediately declare a link down if a hello message is lost once in a while.

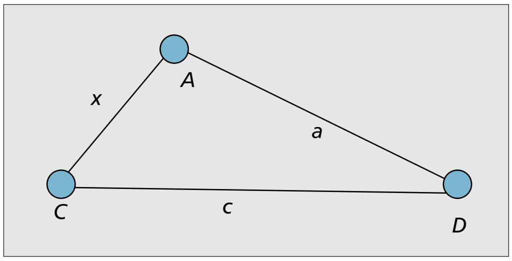

Figure 2. *Selecting the best neighbor A in localized routing schemes.*

Localized Position-Based Routing

We now address the routing task when nodes are equipped with position information. It is widely recognized that sensor networks need geographic location information to perform their duties. We assume that nodes are aware of their own position, the position of their neighbors, and the position of the destination. Enabling nodes with position information is currently an active research area. A survey of position-based routing protocols is given in [4].

We now describe a general simple design principle for describing greedy position-based routing protocols when the destination position is known. Node A, currently holding the packet, will forward it to neighbor B, closer to the destination than itself, which minimizes the *ratio of cost over progress*. The cost measure depends on the assumptions and metrics used, while progress measures the difference in distances to the destination. Several known protocols are special cases of this design.

Let C be the node currently holding the message, D the destination node, A the considered forwarding neighbor, $|CD| = c$, $|AD| = a$, and $|CA| = x$ (Fig. 2). Several localized position-based algorithms are described in [5]. We describe only those that follow the described general design principle, which are also the best performing ones. The *progress* made by forwarding from C to A is $c–a$. The *cost* measure considered is EHC, which depends on distance (which impacts the selection of optimal number u of acknowledgments) and other assumptions made. The optimal number of acknowledgment retransmissions u is approximated as $u \approx 1/p(x)$. If packets are of fixed size and acknowledgments are of the same packet size, the expected hop count is $f(u,x) = 2/(p(x)(1 - (1 - p(x))^u))$ [5]. If other assumptions are made, the corresponding value of $f(u,x)$ needs to be found, the rest of protocol being the same. Node C currently holding the packet will forward it to a neighbor A (closer to the destination than itself) that maximizes the ratio $(c–a)/f(u,x)$ of expected progress and cost for the progress made [5]. If acknowledgment packets are of different size, the described algorithms are still applicable by only changing the corresponding formulas involving acknowledgments.

We now describe another general design paradigm for localized routing protocols. For any given algorithm, an iterative improvement version can be described as follows. Suppose that node C, currently holding the message, selects neighbor A based on a given metric and protocol (e.g., cost/progress ratio). Then an intermediate common neighbor node B (closer to the destination than C, if it exists) is found that minimizes $cost(CB) + cost(BA)$. If $cost(CB) + cost(BA) < cost(CA)$, B becomes the new forwarding neighbor, taking the role of A. This process is iteratively repeated until no improvement is possible. Node C will forward the message to the selected neighbor A, which then applies the same scheme for its own forwarding.

Now consider the model that does not have hop-by-hop acknowledgments. The progress made by forwarding from C to A is $c–a$. This progress is probabilistic. In a nonacknowledged progress routing (*nEPR*) algorithm [6], node C currently holding the message will forward it to a neighbor A (closer to the destination than itself) that maximizes the expected progress, which is the product of the probability of successful delivery $p(x)$ of the message from C to A and the progress made $(c–a)$ by forwarding to A. Therefore, the neighbor A that maximizes $p(x)(c–a)$ is chosen to forward the message. Note that $1/p(x)$ could be considered a measure of cost in this scenario. The iterative improvement can be applied similarly for this scheme, using the product of probabilities instead of the sum of costs in the comparison.

The described routing algorithms and several others are simulated in [3, 5, 6], using a model without collisions from other traffic and with the ideal MAC layer (each link has cost exactly as in the EHC metric). The results clearly show that a shortest path algorithm with the hop count metric performs poorly when it is evaluated with the new metric, since long links with many retransmissions are preferred. The ideal routing protocol is obviously the shortest weighted path algorithm, with EHC as the cost of each link. Our simulations show very competitive performance (closely matching the performance for sufficiently dense networks) of localized routing protocols compared to the ideal shortest weighted path solution, which requires global network knowledge (and therefore unacceptable communication overhead for its maintenance) at each node. We also show that the localized progress-based routing protocol is superior to the localized threshold-based greedy protocol for any selection of threshold. In the latter protocol, the message is sent to the neighbor closest to the destination among neighbors that are at distance $\leq tR$, where t is a parameter threshold. Therefore, when hop count is replaced by EHC, it is not sufficient to merely clarify the transmission radius in order to design optimal localized routing protocol.

To illustrate further our general design guidelines, consider the next extension of physical layer model by adding variable bit rates (which depend on channel properties that in turn depend on distances). The optimality metric then becomes the expected time for sending a packet (instead of the expected hop count) in case of routing with acknowledgments. If we assume that bits are received independent of each other and with equal probabilities, it is relatively straightforward to compute the new cost $f(u,x)$ (the expected time) on each link, and progress-based routing protocol can be still applied. However, if we assume that the corresponding channel modulation and error correction schemes for each bit rate are applied (this changes our assumption about each bit being independently received), the computation of $f(u,x)$ becomes complex, but the routing protocol is otherwise still applicable.

We have so far only described physical-layer-based solutions for greedy position-based routing. Routing with guaranteed delivery for the UDG model and ideal MAC layer was described in [8], which proposed the Greedy-Face-Greedy (GFG) procedure. When a node has no neighbor closer to the destination than itself, it resorts to face recovery mode until a node closer to it is found. The recovery procedure is based on a planar graph locally defined. This procedure can be adapted to the physical layer in a straightforward manner. The edges of the planar graph are normally short ones, thus having relatively high reception probabilities. They are therefore good choices for edge selection. Hence, the recovery mode for physical layer impact routing may proceed in the same way as in the UDG model. Only the greedy mode needs to be changed.

One interesting extension of position based routing described here is beaconless routing, where nodes are not aware of existing neighbors when they want to forward the message. Zorzi [9] proposed to avoid duplicate forwarding in the beaconless routing scheme by applying the RTS/CTS MAC scheme. The current node sends the RTS signal instead of a message, and waits for a node to respond with a CTS signal. The neighbors closer to the destination will respond sooner, based on an appropriately set timeout for the response. If several responses are received, the node selects one that appears to be the best for forwarding and then sends the packet to that neighbor directly. This procedure can be adapted to the physical layer by modifying the criterion for selecting the best forwarding neighbor and appropriate timeout. The timeout can be based on the formulas already described here for selecting the best forwarding neighbor. If a given node announces the request to forward a packet several times, the best forwarding neighbors will receive it, and that neighbor will respond a few times to make sure the response was received and it was selected.

Physical-layer-based recovery schemes and beaconless routing may be combined into a single scheme to produce beaconless routing with guaranteed delivery (of course, this claim is subject to some obvious conditions since the performance is always probabilistic). A simple solution is that the current node C solicits responses from all neighbors and selects the proper one after gaining neighbor knowledge. We are designing a more optimal solution that is not reported here due to space constraints.

CONCLUSIONS

Users of real ad hoc and sensor network equipment will soon observe that the UDG model, although a good initial approximation, is insufficient to produce optimal network layer algorithms for practical applications. For instance,

Our simulations show very competitive performance (closely matching the performance for sufficiently dense networks) of localized routing protocols compared to the ideal shortest weighted path solution, which requires global network knowledge at each node.

An interesting open problem for future research is to consider physical-layer-based routing and broadcasting where nodes may adjust their transmission radii. Expected power consumption may then be considered a primary optimality measure.

the UDG model suggests forwarding the packet to the neighbor closest to the destination to minimize hop count. This may be counterproductive, as signal strength variations will cause low packet reception probabilities at far distance, therefore requiring several retransmissions for the packet and increased expected hop count. We expect that research on physical layer impacts on routing and broadcasting, and other network layer protocols (e.g., geocasting, multicasting, location service) will soon grow significantly.

An interesting open problem for future research is to consider physical-layer-based routing and broadcasting where nodes may adjust their transmission radii. Expected power consumption may then be considered a primary optimality measure. Some preliminary discussions are given in [3].

Further research should address other problems in the design of network layer protocols. For instance, if we consider a more dynamic and realistic channel model, such as multipath fading, the estimated number of packets may suffer from large variance, and the described protocols may need some adjustments. More realistic interference models can be added, and transport layer protocols also need to be adjusted.

ACKNOWLEDGMENTS

The authors greatly appreciate constructive comments received from referees and editors that helped us in preparing this significantly improved and modified version of this article. This research is partially supported by NSERC.

REFERENCES

[1] T. Nadeem and A. Agrawala, "IEEE 802.11 Fragmentation-Aware Energy-Efficient Ad Hoc Routing Protocols," *Proc. 1st IEEE Int'l. Conf. Mobile Ad Hoc and Sensor Systems*, Fort Lauderdale, FL, Oct. 2004.

[2] L. Qin and T. Kunz, "On-demand Routing in MANETs: The Impact of a Realistic Physical Layer Model," *Proc. 2nd Int'l. Conf. Ad Hoc Networks & Wireless*, 2003, pp. 37–48.

[3] I. Stojmenovic *et al.*, "Physical Layer Impact on the Design And Performance of Routing And Broadcasting Protocols in Ad Hoc And Sensor Networks," to appear, *Comp. Commun*.

[4] S. Giordano and I. Stojmenovic, "Position Based Routing in Ad Hoc Networks, A Taxonomy," *Ad Hoc Wireless Networking*, X. Cheng, X. Huang, and D.Z. Du, Eds., Kluwer, 2003.

[5] J. Kuruvila, A. Nayak, and I. Stojmenovic, "Hop Count Optimal Position Based Packet Routing Algorithms for Ad Hoc Wireless Networks with a Realistic Physical Layer," *Proc. 1st IEEE Int'l. Conf. Mobile Ad Hoc and Sensor Systems*, Fort Lauderdale, Oct. 2004; to appear, *IEEE JSAC*.

[6] J. Kuruvila, A. Nayak, and I. Stojmenovic, "Greedy Localized Routing for Maximizing Probability of Delivery in Wireless Ad Hoc Networks with a Realistic Physical Layer," *CD Proc. 1st Int'l. Wksp. Algorithms for Wireless and Mobile Networks*, Mobiquitous, ICST, Boston, MA, Aug. 2004.

[7] D. De Couto et al., "A High-throughput Path Metric for Multi-hop Wireless Routing," *Proc. 9th Int'l. ACM Conf. Mobile Comp. & Net*, San Diego, CA, 2003.

[8] P. Bose *et al.*, "Routing with Guaranteed Delivery in Ad Hoc Wireless Networks," *ACM Wireless Networks*, vol. 7, no. 6, Nov. 2001, pp. 609–16.

[9] M. Zorzi, "A New Contention-based MAC Protocol for Geographic Forwarding in Ad Hoc and Sensor Networks," *IEEE ICC*, Paris, France, 2004.

BIOGRAPHIES

AMIYA NAYAK (anayak@site.uottawa.ca) is an associate professor at SITE, University of Ottawa. He had over 15 years of industrial experience in software engineering in avionics and telecommunication applications prior to joining University of Ottawa in 2002. He has been an adjunct research professor at Carleton University since 1994 and the Canadian Editor and Book Review Editor of *VLSI Design: An International Journal of Custom-Chip Design, Simulation, and Testing* since 1998. His research interests are in the areas of fault-tolerant computing, ad hoc networks, and distributed computing.

IVAN STOJMENOVIC (ivan@site.uottawa.ca) has published over 170 different papers in journals and conferences, edited the *Handbook of Wireless Networks and Mobile Computing* (Wiley, 2002), and co-edited *Mobile Ad Hoc Networking* (IEEE Press, 2004). His current research interests include wireless ad hoc, sensor, and cellular networks. He is currently editor of several journals including *Journal of Multiple-Valued Logic and Soft Computing*, *IEEE Transactions on Parallel and Distributed Systems*, *Parallel Processing Letters*, and *Parallel Algorithms and Applications*. He guest edited recently special issues in several journals including *IEEE Computer Magazine*, *IEEE Network*, and *Wireless Communications and Mobile Computing*.

JOHNSON KURUVILA (kuruvila@site.uottowa.ca) has 15 years of industry experience in research, software development, embedded/distributed system design, and IP networking/wireless protocols. Earlier he was lead engineer at AcceLight Networks and director of routing/signaling platforms at AcceLight Research. He was also manager of IP routing system development at Nortel Networks. He holds degrees in electronics/communication engineering and computer science. His current research interests are in the areas of wireless ad hoc and sensor networks.