Performance Assessment of a Class of Cross-Layer Optimized Protocols for Geographic Routing in WSNs

Andrea Munari and Wolfgang Schott
IBM Research GmbH - Zurich Research Laboratory (ZRL), Zurich, Switzerland.
{una, sct}@zurich.ibm.com

***Abstract*—Geographic routing is a promising solution for data forwarding in power-constrained ad-hoc networks. Such protocols are especially attractive for wireless sensor networks as they can be operated without a central control and are scalable to an arbitrary number of nodes. In this work, we present a class of geographic routing algorithms that enhance energy-efficiency by means of a cross-layer design that incorporates PHY/MAC functions such as on-demand cooperative relaying and leapfrogging. The performance of these protocols is evaluated by means of OMNET++ simulations. Significant improvements are shown in terms of packet delivery ratio, latency and connectivity for different network topologies.**

I. Introduction and Related Work

Wireless sensor networks (WSNs) are typically composed of a large number of nodes that are limited in power and computational capabilities. Since the sensor nodes are often scattered over a wide geographic area, multi-hop transmissions have to be applied to forward the sensed data from a source to a remote destination. Due to the characteristics of WSNs, the employed routing schemes have both to preserve the battery energy of nodes, enhancing network lifetime, and to perform well with an arbitrary number of sensor nodes. *Scalability* and *energy efficiency* are thus the two primary constraints for the design of these algorithms.

A *scalable* solution for data forwarding in WSNs can be obtained by using geographic routing [1]. These protocols assume nodes to be location aware. Under this hypothesis, a route to the sink can be dynamically established while the information is being forwarded. A node that has a packet to deliver transmits a request message to its neighbors containing its own position and the location of the addressee. Nodes that decode this packet start contending as next hop node by setting up a timer which is inversely proportional to their distance to the destination. The node that offers the maximum advancement towards the sink accesses the medium first and replies to the source request, being selected as next hop in a fully distributed way. Once such a channel negotiation has come to an end, data are forwarded to the chosen relay, that iterates the procedure. Since the next-hop node selection procedure only requires a local information exchange between the sensor and its neighbors, the algorithm scales well for networks with a large number of sensors [2]. It is well-known that routing protocols that do not use geographic location information for path selection, such as ad-hoc on demand distance vector routing [3], are not scalable.

Energy efficiency in WSNs can be obtained by designing the network control algorithms with a cross-layer approach [4], jointly defining protocol functions that are assigned to separate layers in the classical communication model. An example of cross-layer optimized geographic routing scheme has been introduced in [5]. This protocol incorporates the medium access control (MAC) concepts of on-demand cooperative relaying [6] and leapfrogging at the network layer in order to improve energy efficiency by taking advantage of the broadcast nature of the wireless medium.

In this paper, we present and discuss a class of cross-layer optimized protocols for routing in energy-constrained WSNs. These algorithms assume location-aware sensor nodes and apply the geographic routing paradigm to achieve scalability. Moreover, the proposed solutions stress the cross-layer approach by both considering information on the instantaneous state of radio links and residual battery energy level of sensor nodes for the path selection algorithm and by incorporating into the routing procedure some MAC layer concepts. The performance of the protocols is assessed by means of OMNET++ [7] network simulations.

The remainder of the paper is organized as follows: Section II presents the proposed cross-layer optimized routing protocols; Section III describes the considered simulation environment and the underlying network assumptions; Section IV discusses the simulation results and Section V concludes the paper with an outlook on future work.

II. A Class of Cross-Layer Optimized Geographic Routing Protocols

In this Section, we present and discuss a class of cross-layer optimized geographic routing algorithms tailored for WSNs, namely Distributed Routing, Cooperative Distributed Routing and Distributed Routing with Cooperative Relaying and Leapfrogging, together with a reference protocol for schemes that rely on static paths, called Shortest Path Routing. All the algorithms assume a CSMA/CA protocol to control the access to the radio channel.

A. Shortest Path Routing (SPR)

Shortest Path Routing follows a greedy approach for data forwarding. The next hop is always the closest node to the destination among the set of neighbors. No distributed route construction procedure is applied, and the path from a source to a sink is uniquely identified by the network topology.

B. Distributed Routing (RoDi)

Geographic routing algorithms typically consider advancement towards the destination as the only criterion to identify the next hop. However, this approach has some important drawbacks. First of all, a minimization of the route length does not necessarily induce a reduction of the overall number of transmissions required to deliver data to the destination. In fact, the best positioned node may experience poor channel

978-1-4244-2644-7/08/$25.00

conditions with the source, and a data transmission addressed to it may fail, triggering a retransmission attempt. Secondly, the lifetime of sensor nodes should be taken into account. If a greedy approach is used, nodes that are closer to the destination are exploited more often for data forwarding, and thus incur a higher energy consumption. This likely results in early battery depletion and may induce holes in the topology with negative effects on the overall packet delivery capabilities.

Starting from these remarks, we propose a geographic routing protocol called Distributed Routing that introduces an enhanced version of the metric used for the distributed selection of the next hop. When a source S has data to deliver to a destination D, it transmits an RTS message with its own location and the location of the destination. A node x that decodes the packet estimates the channel gain $|h_{S,x}|$ with S and computes a metric m_x that takes into account its distance to the destination, $d_{x,D}$, $|h_{S,x}|$, and its remaining battery energy E_x:

$$m_x = E_x \cdot \frac{|h_{S,x}|^2}{d_{x,D}^{\alpha}}\,, \tag{1}$$

where α is the path loss coefficient of the radio channel. This metric increases for nodes with high channel gain to the source (i.e., likely to decode a successive data packet), high remaining energy and which are close to the destination. Once the computation is performed, a backoff timer inversely proportional to m_x is set and a contention procedure is started. During the backoff period, the node listens to the medium. If the sensed power exceeds a given threshold, the node gives up the contention and goes back to the idle state. On the contrary, if the timer expires, the node transmits a CTS frame with its ID, proposing itself as next hop. Notice that this procedure manages to select a relay and to negotiate the channel according to CSMA/CA at once.

C. Cooperative Distributed Routing (RoCoDi)

Let us assume that a node x has been selected as next hop by a geographic routing algorithm and suppose that the subsequent data packet cannot be decoded. In this case, a protocol like RoDi follows the CSMA/CA approach and requires the preceding node to perform another attempt addressed to x. However, such a retransmission is not likely to succeed before the channel conditions that induced the failure change favorably. On the other hand, thanks to the broadcast nature of the wireless medium, other nodes may have decoded the data packet even though it was not intended for them. The cooperative paradigm proposes that one of these nodes immediately retransmits the frame in place of the preceding node. In this way, two copies of the same packet received over independent channels are available at x, that can perform Maximum Ratio Combining, significantly increasing the probability of a successful decoding by taking advantage of spatial diversity [6].

In this work, we present a protocol called RoCoDi that extends RoDi in order to exploit cooperative relaying. Let $\mathcal{R}$ be the set of nodes that decode a data packet sent by the preceding node. If the reception at the next hop x fails, the node transmits a NACK frame asking for a retransmission. A node $r \in \mathcal{R}$ that receives this packet computes a metric m_r and sets up a backoff timer inversely proportional to it. The metric stems from the same principles that lead the next hop selection, and is defined as:

$$m_r = E_r \cdot \frac{|h_{r,x}|^2}{d_{r,x}^{\alpha}}\,, \tag{2}$$

where the channel gain between r and x is estimated from the reception of the NACK frame. During the backoff phase, carrier sensing is performed. If the power on the medium exceeds a given threshold, the node assumes that someone else wins the contention and goes back to idle state. On the contrary, if the timer expires, the node acts as relay by transmitting cached data. At the end of this phase, x either sends out an acknowledgment and proceeds with the forwarding or transmits a further NACK (if data were not decoded or no relay was present), passing the token back to S (that may perform another attempt or trigger a link failure).

D. Distributed Routing with Cooperative Relaying and Leapfrogging (RoCoDiLe)

The broadcast properties of the wireless medium offer, besides the discussed advantages of cooperation, the possibility of opportunistically shorten routes. Let us assume that a cooperative relaying phase has taken place with a node r having performed a data retransmission addressed to x on the behalf of S. Nodes that are closer to the destination than x may happen to decode the relayed data packet even though it was not addressed to them. In this case, if one of such nodes takes over the task of forwarding data in place of x, it is possible to further approach the sink without undergoing additional hops. This mechanism is called *leapfrogging* and aims at reducing the number of transmissions required to deliver the payload, inducing potential gains in terms of delay and energy consumption. These benefits come at the expense of an increased protocol complexity required to coordinate the nodes that take part in the data exchange.

In this work, we discuss a protocol called RoCoDiLe, proposed in [5], that extends RoCoDi by including a distributed algorithm based on carrier sensing to exploit leapfrogging. Let $\mathcal{L}$ be the set of nodes that successfully decode a data packet sent to x by a relay. A node in $\mathcal{L}$ compares its distance from the destination to the distance of x[1] from the destination. If the node, say l, is closer than the current next-hop to D, a carrier sensing contention procedure resembling the ones described for RoDi and RoCoDi is started by setting up a backoff timer inversely proportional to a metric m_l, computed as:

$$m_l = \frac{E_l}{d_{l,D}^{\alpha}}\,. \tag{3}$$

The node that wins the contention transmits a leapfrog request (LPFREQ) message addressed to x. If this packet is correctly received, x replies with an acknowledgment frame (LPFACK) and does not proceed with data forwarding. On the other hand, the leapfrog node takes the role of next hop only upon the reception of an LPFACK in order to avoid flow duplications in the network.

III. SIMULATION ENVIRONMENT

The protocols described in Section II have been evaluated by means of OMNET++ network simulations. The networks were composed by 25 sensor nodes spread inside a 200m ×

[1]The position of the current next-hop could be included in the data packet sent by the relay.

200m area with one source and one destination. Two types of network topologies have been considered: grid and random ones. In the former case, the nodes are disposed to form a 5 $\times$ 5 grid with a distance of 25 meters between any two vertically or horizontally aligned nodes. The source is located on the upper-left corner and the sink on the lower-right corner. In the latter case, the positions of the source and the destination are kept unchanged, but all the other nodes are uniformly and independently distributed over the area. In our simulations, we have only considered random topologies which are connected, where at least one route from the source to the sink can be formed with neighboring nodes.[2]

The wireless environment is subject to path loss with exponent $\alpha = 3.5$ and correlated Rayleigh fading. Transmissions are performed at an information rate $R = 2$ bit/s/Hz. Defining the instantaneous channel capacity as $C(t)$ [6], an outage event occurs if $C(t) < R$. In such a condition, the packet reception is assumed erroneous, otherwise it is correct.

All the nodes in the network have an initial battery energy equal to 30 J. The power consumption in reception mode has been set to 30 mW, whereas the transmission power has been varied in order to obtain a target SNR.

The source injects traffic at a low rate of 0.4 packets per second. Signaling packets are 16 bytes long, while data packets are composed of 128 bytes. The behavior of the protocols has been assessed by averaging the results obtained performing 20 simulations that lasted 1000 seconds each. In the random topology scenario, each run corresponded to a different placement for the nodes.

IV. Simulation Results

A. Grid Topology

Fig. 1 shows the metric Packet Delivery Ratio (PDR) as a function of average SNR for the considered protocols. The PDR is defined as the ratio of the number of packets successfully received at the destination to the number of packets injected in the network by the source. The average SNR is experienced at a receiver, whose distance from the transmitter equals the length of one hop over the main diagonal.

For high values of SNR, the PDR tends to one for all the protocols as expected. In the low-SNR regime, SPR performs very poorly because it relies on a single static path and thus the probability of reaching the destination is very low. Geographic routing significantly mitigates this problem. With this strategy, if the link with the neighbor which is closest to the destination is not good enough to ensure a successful data transfer, other forwarding opportunities are considered and exploited, enhancing the chances of delivering data to the destination. If RoDi is applied, a PDR gain up to 40% for SNR values between 10 and 15 dB is achieved. Further improvements are obtained if cooperative relaying is employed as can be seen by comparing the curves for RoDi and RoCoDi. For low-SNR values, the two protocols perform similarly. In these conditions, a link between two nodes is likely to fail during the contention phase and therefore cooperation is seldom used. As transmission power is increased, relaying starts to show its influence and gains up to 10% are achieved for mid-SNRs. The cooperative advantage is twofold: not only packet error rates are reduced because of spatial diversity, but also ARQ phases are triggered more often as the retransmission request is more likely to be correctly decoded, being addressed to a set of nodes and not only to the original data transmitter. As far as RoCoDiLe is concerned, some limited improvements over RoCoDi are obtained. This result is not surprising, as leapfrogging (unlike cooperation) is not meant to increase the robustness of established links but rather to decrease the latency. The slight increase in PDR is due to the reduction of the average route length offered by RoCoDiLe (refer to the discussion of Fig. 2), as shorter paths are statistically less likely to induce packet losses.

Fig. 2 presents the latency per packet delivery for the different protocols, defined as the average time needed to successfully deliver data from the source to the sink. The metric is depicted as a function of average SNR and has been normalized to the duration of a data packet transmission between two nodes. First of all, we notice that the average latency for SPR decreases as transmission power is increased, because of the reduced number of retransmissions required per successful packet delivery. In the low-SNR region, RoDi and RoCoDi are able to deliver to the sink many more packets than SPR (see Fig. 1) but with a higher average latency. In these conditions, when a node sends out an RTS packet, the probability that some of its neighbors do not correctly decode it is not negligible. Therefore, the choice of the next hop is

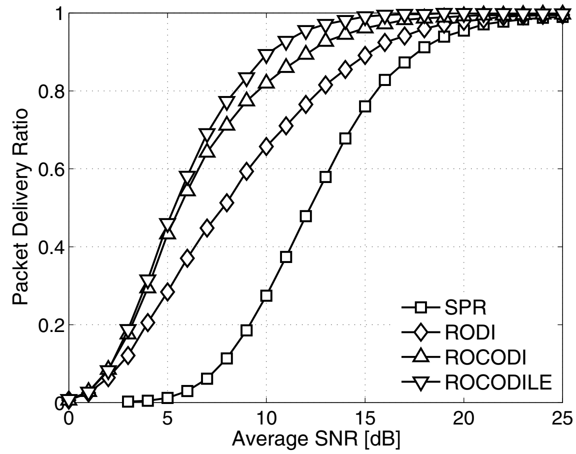

Fig. 1: Packet Delivery Ratio vs SNR, grid topology

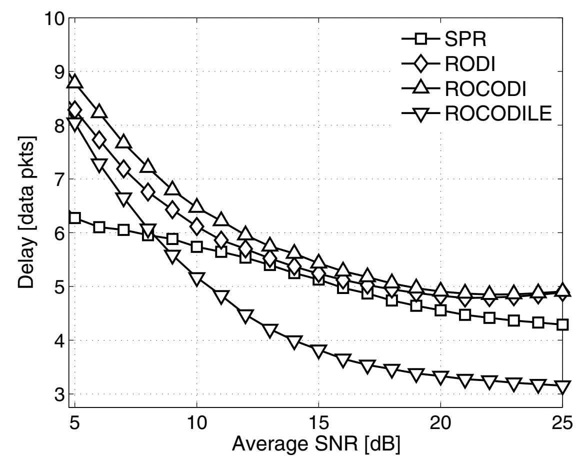

Fig. 2: Delay vs SNR, grid topology

[2]We consider as neighbor nodes whose relative distance is at most equal to the length of one hop over the main diagonal of the grid topology.

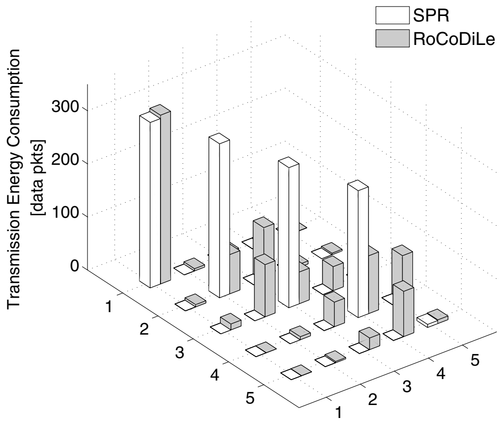

Fig. 3: Transmission energy consumption, grid topology. Source node is located at (1,1) while destination node is at (5,5)

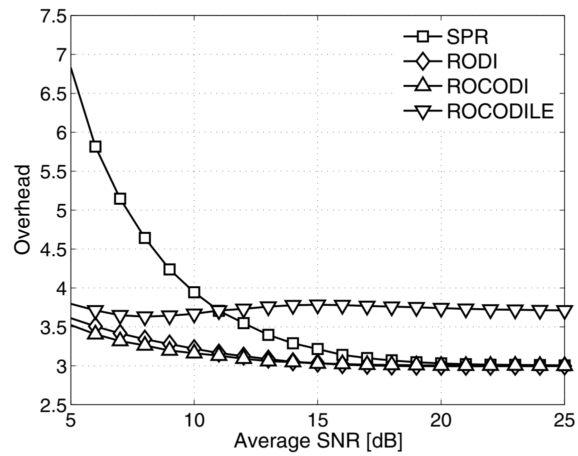

Fig. 4: Average overhead vs SNR, grid topology

performed among a subset of the potential candidates, and the possibility of picking up routes different from the shortest one increases with a subsequent effect on latency. As SNR raises, this effect is mitigated and the average delay lowers as well. For low and medium values of transmission power, RoCoDi performs worse than RoDi. Cooperation increases the probability of delivering to the sink packets that incur a deep fade. However, such data flows experience a longer delay due to (successful) retransmissions and contribute to increase the average latency metric for RoCoDi. As far as RoCoDiLe is concerned, important advantages over the other geographic routing schemes are achieved in terms of average delivery time thanks to the opportunistic bypassing of some hops. The benefits induced by shorter paths overcome the latency introduced by the channel negotiation necessary to set up the leapfrogging procedure (i.e., LPFREQ and LPFACK) and RoCoDiLe is able to outperform even SPR for sufficiently high levels of transmission power. In the low-SNR region, the protocol performs slightly worse than SPR as leapfrogging is seldom triggered due to poor reception conditions.

Fig. 3 depicts the transmission energy consumption in the network. We compare SPR and RoCoDiLe operating with the same transmission power (SNR = 15 dB). The x and y coordinates identify a node in the network topology, the source being located at (1,1) and the destination being at (5,5). The z coordinate represents battery consumption due to transmissions and for the sake of clarity it has been normalized to the energy required to perform a single transmission of a data packet. First of all, the plot highlights that the source experiences the highest consumption regardless the protocol implemented. This is due to the fact that all the packets are generated there and therefore the node is involved in each communication. On the contrary, the sink shows a low power consumption as it only transmits control packets. Moreover, energy usage degrades along the main diagonal: the closer a node is to the sink, the lower the probability that it is reached by a frame being forwarded, as more hops need to be successfully performed. SPR only consumes the battery of nodes along the shortest path. On the other hand, the usage of geographic routing together with cooperative relaying and leapfrogging involves many more nodes in data flows. In this way energy consumption is much more distributed over the network and the average battery level is significantly higher for RoCoDiLe. In conclusion, not only the class of protocols that we propose is able to increase the efficiency of routing in terms of PDR and delay, but also network lifetime is significantly improved.

Finally, we consider the average overhead per data transmission induced by the protocols, depicted against average SNR in Fig. 4. The metric is defined as the ratio of the total number of signaling packets sent in the network to the number of sent data packets (either as first transmissions, ARQ phases or relaying). The overhead for protocols that do not employ leapfrogging asymptotically tends to 3, as only the CSMA/CA negotiation packets are required (RTS, CTS, ACK). For low SNR, SPR presents a significant increase of the metric due to the higher number of ARQ phases undergone. The same reasoning explains why RoDi performs slightly worse than its cooperative version (RoCoDi). The higher overhead that characterizes RoCoDiLe is required in order to coordinate nodes that take part in the leapfrogging procedure and to avoid flow duplications and congestion in the network. For low transmission powers this mechanism is not often used and thus the additional signaling is limited. Leapfrogging is exploited more often in the mid-SNR region, as confirmed by both Fig. 4 and Fig. 1. Finally, for high values of transmission power the overhead introduced by RoCoDiLe slightly lowers as it is counterbalanced by the reduction in the number of required retransmissions (resembling the trends for RoDi and RoCoDi).

B. Random Topologies

Fig. 5 depicts the PDR vs average SNR for the considered protocols in random topology scenarios. Generally speaking, the trends discussed for the grid topology are confirmed. However, for high values of the transmission power, the metric does not converge to 1. For SPR, this is due the greedy approach followed to select the next hop. As discussed, in our simulations we have considered only topologies where at least one path to the sink exists. However, this assumption does not guarantee that one of such paths can be found using maximum geographic advancement as the only criterion. As a result, SPR may not be able to deliver packets in some topologies, regardless the quality of the link, and therefore the PDR is lowered. Geographic routing protocols significantly mitigate this effect by exploiting the distributed selection of the next

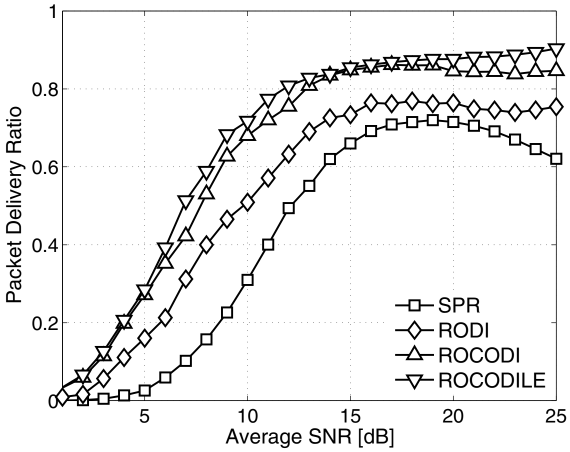

Fig. 5: Packet Delivery Ratio vs SNR, random topology

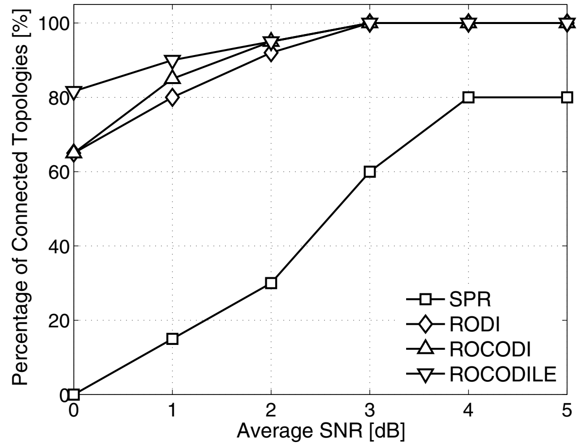

Fig. 6: Connectivity, 25 nodes network

hop, but still not all the packets are delivered successfully to the sink as some deadlock situations may occur (i.e., the dynamic selection of the next hop may pick up a node that does not have any connection to the final destination). It is also interesting to observe that the curve of SPR decreases for extremely high values of SNR. This trend is due to the high transmission energy consumption that characterizes this activity region and that may lead some nodes to run out of battery within the simulation interval. When such conditions occur, SPR is not able to reach the sink anymore, as only one route is known and exploited and PDR tends to drop. Such a trend does not affect the other protocols. The reason is twofold. On one hand, the better distribution of energy consumption in the network introduced by these schemes reduces the probability of node deaths. On the other hand, even if a node runs out of battery, the distributed selection of the next hop is able to dynamically switch to alternative paths, thus preserving high values for the delivery ratio.

Finally, we analyze the connectivity properties offered by the different protocols. We say that a network topology is connected using a specific routing scheme if at least one packet is delivered to the sink within the simulation interval. Fig. 6 depicts the percentage of connected topologies in the low-SNR region. The metric is defined as the ratio of the number of connected scenarios using a given routing algorithm to the number of simulated scenarios. Let us first consider SPR. For sufficiently large values of transmission power, the percentage of connected scenarios stabilizes to 80%. This shows that a greedy approach that considers geographic advancement as the only criterion for the next hop selection fails in determining a route in about one-fifth of the topologies, regardless the link quality. On the other hand, for low-SNRs the number of non-connected scenarios significantly increases due to channel impairments that make it more difficult to reach the sink even if one greedy path exists. As far as geographic routing protocols are concerned, for SNR larger than 3 dB all the scenarios are connected. This shows that the distributed choice of the next hop significantly increases the reliability of routing by taking into account paths that differ from the shortest one. The robustness to channel impairments offered by cooperative relaying and leapfrogging further stresses the gains over SPR in the very low-SNR region, with improvements up to 60%.

V. Conclusions and Future Work

In this paper, we have presented a class of geographic routing protocols for WSNs that aim at improving robustness and energy efficiency. This objective is pursued by taking into account cross-layer information when computing the metric used to dynamically select the next hop and by incorporating the MAC concepts of cooperative relaying and leapfrogging. The protocols have been compared to a non-geographic routing scheme that relies on the shortest path to deliver data. Significant advantages have been shown in terms of packet delivery ratio, latency and energy consumption without significantly increasing the protocol overhead. Future research directions include a refinement of the proposed schemes as well as the study of the impact of mobility on geographic routing.

Acknowledgments

The authors would like to thank R. Doss and P. Coronel for their contributions. This paper describes work undertaken in the context of the e-SENSE project, "Capturing Ambient Intelligence for Mobile Communications through Wireless Sensor Networks" (www.ist-e-SENSE.org). e-SENSE is an Integrated Project (IP) supported by the European 6th Framework Programme, contract number: 027227.

References

[1] M. Zorzi and R. R. Rao, "Geographic Random Forwarding (GeRaF) for Ad Hoc and Sensor Networks: Multihop Performance," *IEEE Transactions on Mobile Computing*, vol. 2, no. 4, pp. 337–348, 2003.

[2] D. Estrin, R. Govindan, J. Heidmann, and S. Kumar, "Next Century Challenges: Scalable Coordination in Sensor Networks," in *IEEE/ACM Mobicom*, Seattle, WA, US, 1999, pp. 263–270.

[3] C. Perkins and E. Royer, "Ad-hoc On Demand Distance Vector Routing," in *IEEE WMCSA*, New Orleans, LA, US, Feb. 1999.

[4] Z. Haas, "Design Methodology for Adaptive and Multimedia Networks," *IEEE Communication Magazine*, vol. 39, no. 11, pp. 106–107, Nov. 2001.

[5] P. Coronel, R. Doss, and W. Schott, "Geographic Routing with Cooperative Relaying and Leapfrogging in Wireless Sensor Networks," in *IEEE GlobeCom*, Washington DC, US, Nov. 26-30 2007.

[6] D. T. J. Laneman and G. Wornell, "Cooperative Diversity in Wireless Networks: Efficient Protocols and Outage Behavior," *IEEE Transactions on Information Theory*, vol. 50, no. 10, pp. 3062–3080, Dec. 2004.

[7] A. Varga, "OMNET++," in *http://www.omnetpp.org*, 2001.