This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE Globecom 2011 proceedings.

Bypassing Holes in Sensor Networks: Load-balance VS. Latency

Goce Trajcevski*,1, Fan Zhou*, Roberto Tamassia†,2, Besim Avci*,1
Peter Scheuermann*,1 and Ashfaq Khokhar‡,3

*Department of Electrical Engineering and Computer Science
Northwestern University
Email: goce,fanz1,besima1,peters@eecs.northwestern.edu
†Department of Computer Science
Brown University
Email: rt@cs.brown.edu
‡Department of Electrical and Computer Engineering
University of Illinois at Chicago
Email: ashfaq@uic.edu

***Abstract*—This work addresses the problem of geographic routing in the presence of holes or voids in wireless sensor networks. We postulate that, once the boundary of the hole has been established, relying on the existing algorithms for bypassing it may cause severe depletion of the energy reserves among the nodes at (or near) that boundary. This, in turn, may soon render some of those nodes useless for any routing (and/or sensing) purposes, thereby effectively enlarging the size of the pre-existing hole. To extend the lifetime of the nodes along the boundary of a given hole, we propose two heuristic approaches which aim at relieving some of the routing load of the boundary nodes. Towards that, our approaches propose that some of the routes that would otherwise need to bypass the hole along the boundary, should instead start to deviate from their original path further from the hole. Our experiments demonstrate that the proposed approaches not only increase the lifetime of the nodes along the boundary of a given hole, but also yield a more uniform depletion of the energy reserves in its vicinity.**

I. INTRODUCTION

One of the sources of disrupting the expected quality of service in Wireless Sensor Networks' (WSN) settings is the occurrence of *holes* – which is, regions void of operational nodes for sensing and/or routing purposes [1]. Such voids can originate as early as the deployment stage of the network, however, they may also be generated during its continuous operation, either because a significant amount of nodes have depleted their battery-reserves, or because of some environmental phenomenon (e.g., a forrest fire).

The problem of detecting and bypassing holes in WSNs has spurred a significant amount of research results in the past few years [2], [8]–[10], [13], [15], [20], [21]. The settings in the existing works range from purely topological/connectivity relationships (e.g., [10]), to assuming some type of a geographic knowledge regarding the locations of the nodes [3], [16]. When it comes to forwarding a packet from a given *source* to a given *sink* node, typical approaches from the latter settings [9], [10], [13], [15] assume a combination of:

- *greedy forwarding*, from the source towards the sink along the shortest path; and
- *around-perimeter forwarding* (a la' face-routing [14]), in order to bypass the hole that has been encountered along the route towards the sink.

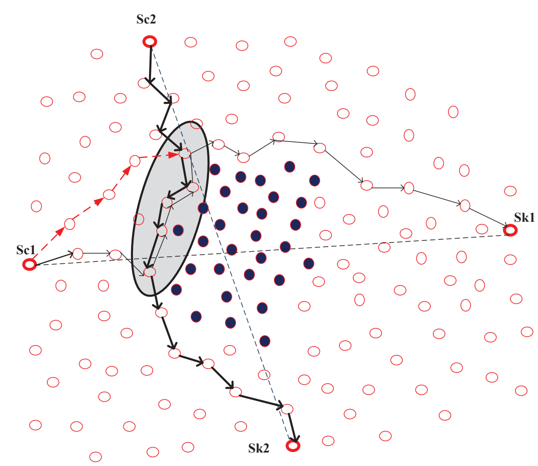

Fig. 1. Motivational Scenario: two routes bypassing the same hole

While the problem of bypassing a given hole using the nodes around its boundary has been investigated in the literature, at the heart of the motivation for this work is the following observation: In scenarios in which there are multiple *(source, sink)* pairs whose routes rely on some (e.g., GPSR-like [13]) protocol to bypass a given hole, their combined-effect may cause the operational nodes along the boundary of the hole to be depleted faster than their neighbors (and the rest of the network). As an additional effect, this could cause a more rapid "expansion" of the hole itself. In addition, routing a particular packet around the boundary of the hole imposes an extra latency on that packet's delivery to the sink. Namely,

(1) Research Supported by the NSF-CNS-0910952
(2) Research Supported by the NSF-CCF-0830149
(3) Research Supported by the NSF-CNS-0910988

if one could, route the packets directly along the tangent from the source to the boundary of the hole, and vice-versa for the sink, the source-to-sink delay will be smaller than the one incurred via greedy geographic algorithms combined with face routing. This is a straightforward consequence of the triangular inequality.

An illustration is presented in Figure 1 which shows two *(source, sink)* pairs, (Sc_1, Sk_1) and (Sc_2, Sk_2) that need to bypass a given hole – depicted as a region containing dark circles (dead nodes). The shaded ellipse points the sequence of nodes that will be used for routing twice, once for each *(source, sink)* pair, thereby doubling the depletion of their energy resources compared to the rest of the nodes along the hole's boundary, not to mention the rest of the nodes in the respective neighborhoods. Observe that routing along the tangent to the hole starting from Sc_1 reduces the hop-count by 3. Both effects will be further amplified in applications in which long-running periodic transmission is required for sampled values from source towards the sink.

The main contributions of this work are methods for compact representation of the hole-boundary and two heuristic-based routing protocols: RCF (Ring-Constraint Forwarding) and WCF (Wedge-Constraint Forwarding) – aiming at balancing the load vs. latency trade-offs. Our experiments demonstrate that each of the proposed heuristics prolongs the lifetime of the boundary nodes and improves the uniformity of the rate of energy-depletion in the network – expressed as a standard deviation of the energy reserves over time.

The rest of this paper is structured as follows. In Section II we present the basic definitions and recollect the necessary background. Section III presents the details of our proposed approaches, and the experimental observations regarding their benefits are given in Section IV. In Section V, we position our work with respect to the relevant literature, summarize the paper and outline directions for future work.

II. Preliminaries

We now introduce the basic terminology and definitions used in the rest of this paper, along with the methodology to generate hole representation used in our heuristics (cf. Section 3).

We assume a dense WSN with N nodes, where the sensor nodes $\{sn_1, sn_2, \ldots sn_N\}$ are randomly deployed over the geographic area of interest. Each node sn_i has a unique, fixed physical location represented as a coordinate in a given reference system. Nodes are assumed to have the capability of determining their locations at run-time, either by means of a location hardware, such as a GPS device, or by implementing a location discovery algorithm [7], [17]. Each node sn_i has information about the position and state of its one-hop neighbors $NB(sn_i)$.

Given a set of point-locations $P = \{P_1, P_2, \ldots, P_M\}$ in the Euclidian 2D space, its *convex hull* $CH(P)$ is defined as the smallest convex polygon whose vertices are $V_{CH(P)} \subseteq P$ and all the other points from $P \setminus V_{CH(P)}$ are inside $CH(P)$ [6]. Using the existing techniques (cf. [8] for determining the nodes that are on the boundary of a given hole, we can readily apply the techniques proposed in [20] to determine the convex hull of the boundary nodes. Assuming that the hole's boundary is determined by M nodes, the complexity of the convex hull algorithm will be bound by $O(M \log M)$, with the worst-case of $O(M)$ messages overhead for the boundary nodes.

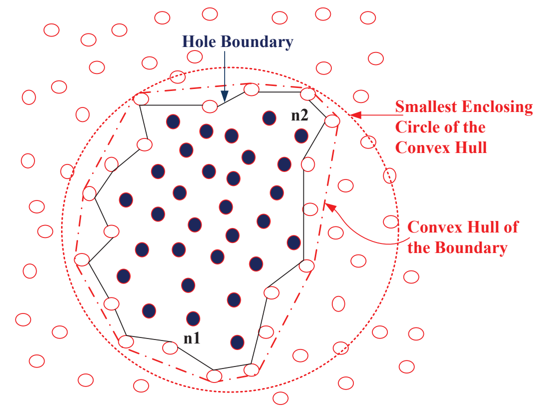

Fig. 2. Convex hull and Smallest Enclosing Circle of a Hole

Let $M_H = |V_{CH(P)}|$ denote the cardinality of the set of vertices in the convex hull of P ($M_H \leq M$). The *smallest enclosing circle* of $CH(P)$ is the circle with the minimal radius which contains every vertex from $V_{CH(P)}$ and, due to the convexity of $CH(P)$ also contains every point from the edges and the interior of $CH(P)$. The construction of the smallest enclosing circle can be performed in $O(M_h \log M_H)$ time complexity ($O(M_H)$ if randomized algorithms are used) [6], with an overhead of $O(M)$ messages for the nodes on the convex hull – since the vertices of the convex hull of the hole need not be within each other's 1-hop communication range.

The concepts that we introduced so far are illustrated in Figure 2. We note that applying a divide-and-conquer strategy, in a manner presented in [20], could yield savings in the messages overheads for both convex hull and smallest enclosing circle computations, however, the upper-bound of $O(M)$ is correct.

III. Hole Expansion and Bypassing

We now present the details of the two routing protocols that aim to balance the trade-offs between two complementary desiderata: reducing the latency of the sink-to-source communication, and the load balancing among the nodes in the network.

A. Ring-Constrained Forwarding (RCF)

The first heuristic that we propose consists of four distinct phases:

Phase I (P_I): Initially, the source S_c starts forwarding along the shortest path towards the sink S_k, e.g., following the $\overline{S_cS_k}$ line segment in a TBF-like manner [18]. If no hole is encountered to be bypassed, the protocol completes the transmission within P_I.

Phase II (P_{II}): **IF** a hole is encountered, the first node along the boundary of its smallest enclosing circle (cf. Section 2) initiates a return-message towards Sc. The content of the message includes the *location of the center* C_c, and the *radius* r_C of the hole's smallest enclosing circle.

Phase III (P_{III}): Whenever P_{II} takes place, it is followed by P_{III}, during which:

1) The source S_c selects the value of d – the width of the *ring* centered at C_c and with inner-radius r_C.
2) It adds the three parameters (C_c, r_C, d) to the packet-content, which will be subsequently carried along the hops.
3) Every subsequent routing node randomly selects as a next-hop one of its neighbors that is:
   – closer towards S_k than itself;
   – is within the zone bounded by the tangents to the inner and outer circles of the ring, as constructed both from S_c and S_k, and the portion of the ring bounded by the circular arcs in-between the tangents.

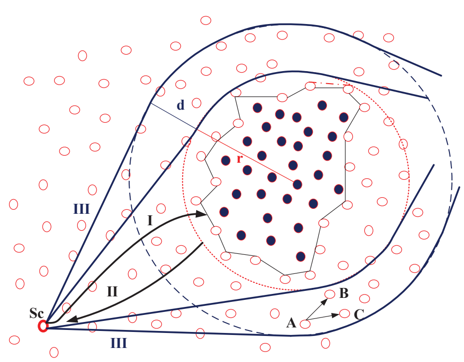

Fig. 3. Back-propagating Hole's Enclosing Circle and Ring Construction

The main ideas behind the RCF protocol are illustrated in Figure 3. As shown, after encountering the hole during *PhI* and having that information propagated back to the source *(PhII)*, S_c selects the value for the width of the ring (d). Note how the sensor node denoted as A in Figure 3 can select any of its neighbors B or C, because they both satisfy the criteria of P_{III}*.3*

B. Wedge-Constrained Forwarding (WCF)

While our first heuristic RCF aims at allowing smaller variations within some bounds from the shortest path around the (smallest enclosing circle of) the hole, the second heuristic that we considered offers a wider choice of selecting next-hop routing nodes.

The first two phases of the WCF, denoted P'_I and P'_{II} are exactly the same as the corresponding ones in the RCF protocol. For that matter, even P'_{III}.1 and P'_{III}.2 are the same with the corresponding parts of RCF. The main difference is in the forwarding policy part of the third phase:

Phase III' (P'_{III}):

1) (same as P_{III}.1)
2) (same as P_{III}.2)
3) Every subsequent routing node randomly selects as a next-hop one of its neighbors that is:
   – closer towards S_k than itself;
   – is within the zone (wedge) bounded by the *tangents to the outer circle*, as constructed both from S_c and S_k, and the portion of the *outer boundary of the ring* in-between the tangents.

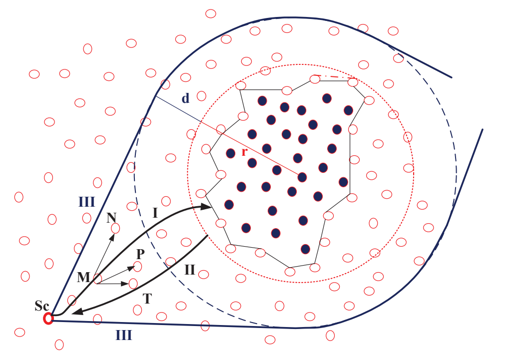

Fig. 4. Bounding Wedge

Figure 4 illustrates the specifics of the WCF heuristics. As shown, the node M can select any of its neighbors N, P, or T – exemplifying the basic trade-offs between RCF and WCF:

- WCF offers wider flexibility of selections, thereby providing better load-balancing;
- RCF, on the other hand, restricts the amount of next-hop selections, however, it provides a smaller latency of the transmission between S_c and S_k.

IV. Experimental Evaluation

We compare our proposed hole-bypassing routing heuristics *RCF* and *WCF* with the Face routing algorithm [4] in several settings.

The experiments were performed on the open source simulator for WSN, SIDnet-SWANS [11]. Each run simulates 1200 homogeneous sensor nodes configured as: (1) 20kbps radio data rate on the MAC802.15.4 protocol; (2) 5 seconds idle-to-sleep interval of inactive nodes to preserve battery power, and 2-seconds interval of data transmission of source node; (3) power consumption characteristics meet the specifications of Mica2 Motes; (4) fully-charged battery fully-charged with initial capacity 25mAh. We evaluate over the lifetime of sensor nodes, load balancing, and the communication time. Each setting was tested for (1) 2 node densities ($\lambda \in \{10, 24\}$ average neighbors per node); (2) 2 hole expanding factors (d = $0.25 r_C$ and $0.75 r_C$); (3) 3 different hole sizes ($Hole$ = 1%, 5% and 15% of field area); (4) Random choice between 1 up to 4 *(source, sink)* pairs. The holes were generated using pentagons and hexagons which were "deformed" by moving the mid-point of the edges towards the interior of the polygon for a randomly selected factor of 20%-60% of the edge's length, and then randomly perturbing the initial location of the vertices within a disk of size 5% of the edges' length. In the sequel, we present the averaged observations of all the runs.

We measure the quality of the load balancing by the standard deviation of the distribution of the energy consumption. As shown in Figure 5, when the hole size is small, e.g., 1% of field area, the effects of the three approaches are very similar (Figure 5(a)). Also, the standard deviation will increase with hole size because more relay nodes are required to bypass

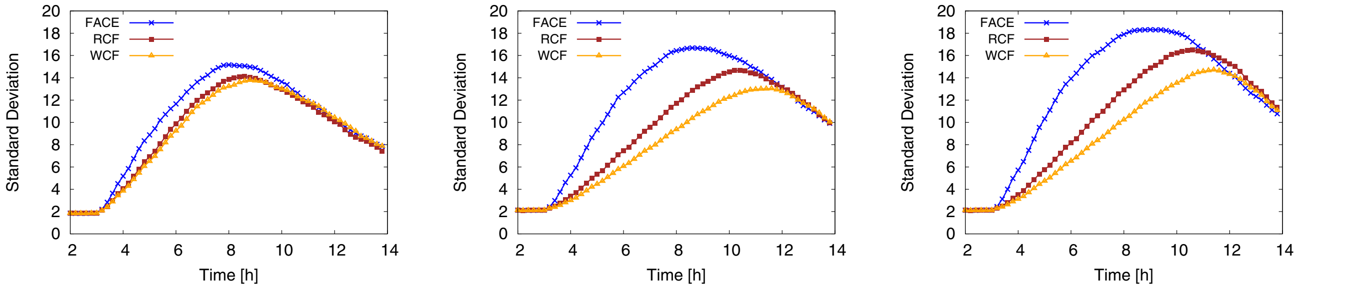

(a) $Hole$ = 1% of field area (b) $Hole$ = 5% of field area (c) $Hole$ = 15% of field area

Fig. 5. Impact of hole size on work-load balance.

the coverage hole. However, face routing is more sensitive to the hole size than *RCF* and *WCF* – due to the lack of path diversity. While the length of routing paths in *RCF* and *WCF* increases with the hole size, more available routing paths can be will reduce the standard deviation of the energy consumption

On average, *RCF* and *WCF* achieve 11% and 28%, respectively, improvement over Face routing in terms of load balancing when hole size metrics are 5% and 15%. Note that the performance gain of *RCF* and *WCF* over Face routing "stabilizes" as the hole size increases to some level (15% in figure 5(c)). This is because the expanding area, decided by factor d, increases with the hole size too, which leads to longer routing paths for both *RCF* and *WCF*.

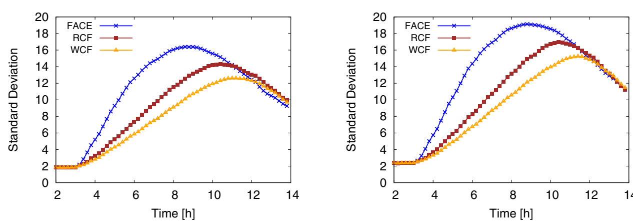

(a) $\lambda = 12$ (b) $\lambda = 24$

Fig. 6. Impact of λ on work-load balance.

Figure 6 shows for $\lambda = 10$, *RCF* and *WCF* achieve improvement (13% and 24%) in load balancing over Face routing. For $\lambda = 24$, the improvement decrease to 11% (*RCF*) and 20% (*WCF*). Note that in both experiments, *WCF* outperforms *RCF*, due to employing more available routing paths.

Figure 7 shows the impact of expanding factor d on the performance of routing schemes *RCF* and *WCF* (for convenience of comparison, we also plot Face routing). As illustrated, increasing d affects the performance of both *RCF* and *WCF* in similar manners. First, increasing d explores more allowable paths for both *RCF* and *WCF* which may utilize a larger fraction of the nodes to share the communication costs. However, increasing d, especially for larger hole size, also increases the paths-length of the *RCF* and *WCF* routing schemes, which may incur more energy consumption, and consequently compensate for the load balancing gain of exploring more available routing paths.

Figure 8(a) shows the life time of nodes, averaged by all parameter settings, where the time is measured based on three

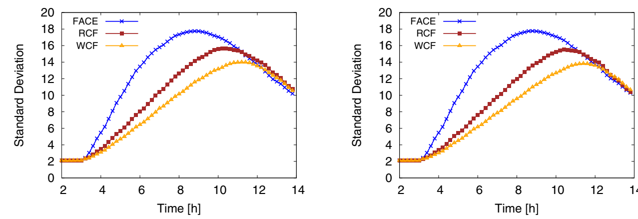

(a) d=0.25 (b) d=0.75

Fig. 7. Impact of factor d on work-load balance

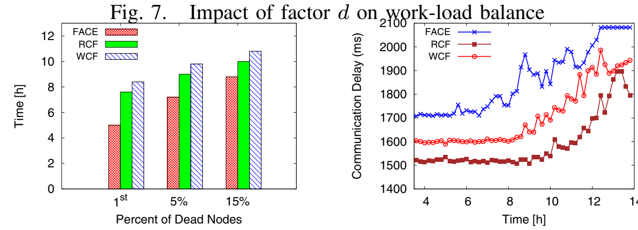

(a) Lifetime (b) Latency

Fig. 8. Trade-offs

different "policies": (1) first dead node; (2) 5% dead nodes; and (3) 10% dead nodes. For a lifetime metric of 15% of dead nodes, *RCF* and *WCF* achieve 1.2 hour (14%) and 2.0 hour (23 %) of additional lifetime than Face routing. When the lifetime metric is reduced to the first dead node, the improvements are even higher (2.6 hour (52%) and 3.4 hour (68 %), respectively). This proves the effectiveness of *RCF* and *WCF* on conserving the energy and extending the lifetime of nodes. Again, *WCF*'s effect on load balancing turn out to be superior, now in terms of lifetime.

Communication latencies are compared in terms of the time to transmit a packet from S_c to S_k. As illustrated in figure 8(b), *RCF* and *WCF* yield 11% and 6% improvements when compared to Face routing. *RCF* performs better than *WCF* due to the use of almost-shortest paths employed to route packets. The latency of each approach increases as the time evolves, because the nodes around the coverage hole die faster, thereby expanding the hole itself, which leads to longer path lengths overall. Figure 8 illustrates the trade-off that we mentioned – while *WCF* achieves better load-balancing (in terms of lifetime), *RCF* incurs smaller latency.

Finally, Figure 9 depicts the total energy dissipation, averaged over all parameters settings. As shown, the proposed routing approaches perform slightly better than Face routing, and *RCF* outperforms *WCF* because it employs fewer hops per routing packet.

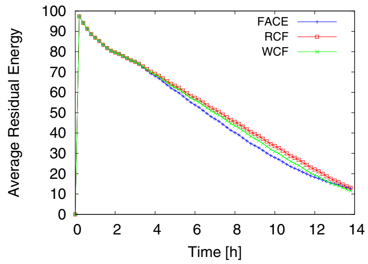

Fig. 9. Average Residual Network Energy

V. Related Work and Conclusions

Problems related to detecting, representing and bypassing holes in WSN settings have been investigated since the emergence of the field (see [1] for a survey). Relying on homology, the problem of coverage and hole detection was addressed in [12] and, more recently, distributed algorithms for finding and patching connectivity holes were presented in [23]. Distributed approaches for locating and bypassing holes based on geometric techniques were used in [8] and in [10] holes' detection methods based on topological properties was presented. Distributed generation of convex hulls was used for coverage holes' detection in [20], and adaptation methods for the networks' operation in presence of holes were suggested in [2]. One immediate consequence of holes is that they need to be considered (and circumvented) for routing purposes. Several geographical protocols have been proposed towards that end [4], [9], [14], [15], [21]. Theoretical analysis for the energy-holes properties in circular multihop networks is presented in [22].

In this work, we specifically addressed the problem of lifetime of the nodes along the boundary of a communication hole in WSNs. Motivated by the observation that those nodes are likely to be more used during bypassing on behalf of queries from different *(sink, source)* pairs, we proposed two heuristics that are trading off the load balancing among the nodes and the communication latency. We proposed to represent the hole compactly via the smallest enclosing circle of the convex hull of the boundary nodes. Upon notification that the shortest geographic path towards the sink contains a hole, and based on its description, the source node decides which routing policy to use (RCF or WCF). Our experiments demonstrated the both of the proposed approaches prolong the lifetime of the nodes along the nodes boundary, when compared against the baseline approach from [4]; RCF providing shorter delay for packets delivery, while WCF providing a better load balancing.

There are two immediate extensions to our work that we plan to pursue in the future. Firstly, we need to take the energy distribution into account (i.e., in this work, we assumed that it is uniformly distributed) [22]. Secondly, we note that in some cases (e.g., polygons that are ”skewed” along a particular dimension) the smallest enclosing circle may introduce a lot of ”dead-space”, in the sense that many nodes near the hole's boundary will not be used at all. We plan to leverage on [5] for more compact representation of holes.

Another extension is along the lines of exploring whether better-controlled multipath routing (e.g., field-based one [19]) could improve the benefits of WCF/RCF.

References

[1] Nadeem Ahmed, Salil S. Kanhere, and Sanjay Jha. The holes problem in wireless sensor networks: a survey. *Mobile Computing and Communications Review*, 9(2):4–18, 2005.

[2] Nadeem Ahmed, Salil S. Kanhere, and Sanjay Jha. Hole healing energy aware algorithms for wireless sensor networks. *International Journal on Computer Theory and Engineering*, 2(5):14–20, 2010.

[3] Amitabh Basu, Jie Gao, Joseph S. B. Mitchell, and Girishkumar Sabhnani. Distributed localization using noisy distance and angle information. In *MobiHoc*, pages 262–273, 2006.

[4] P. Bose, P. Morin, I.Stojmenović, and J. Urrutia. Routing with guaranteed delivery in ad hoc wireless networks. *Wireless Networks*, 7(6), 2001.

[5] C. Buragohain, S. Gandhi, J. Hershberger, and S. Suri. Contour approximation in sensor networks. In *DCOSS*, 2006.

[6] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational geometry: algorithms and applications*. Springer-Verlag New York, Inc., 2001.

[7] Lei Fang, Wenliang Du, and Peng Ning. A beacon-less location discovery scheme for wireless sensor networks. In *INFOCOM*, pages 161–171, 2005.

[8] Qing Fang, Jie Gao, and Leonidas J. Guibas. Locating and bypassing holes in sensor networks. *MONET*, 11(2):187–200, 2006.

[9] Hannes Frey and Ivan Stojmenovic. On delivery guarantees of face and combined greedy-face routing in ad hoc and sensor networks. In *MOBICOM*, pages 390–401, 2006.

[10] Stefan Funke. Topological hole detection in wireless sensor networks and its applications. In *DIALM-POMC*, pages 44–53, 2005.

[11] Oliviu Ghica, Goce Trajcevski, Peter Scheuermann, Zachary Bischoff, and Nikolay Valtchanov. Sidnet-swans: A simulator and integrated development platform for sensor networks applications. In *SenSys*, pages 385–386, 2008.

[12] R. Ghrist and A. Muhammad. Coverage and hole-detection in sensor networks via homology. In *IPSN*, 2005.

[13] B. Karp and H. T. Kung. GPSR: greedy perimeter stateless routing for wireless networks. In *Mobile Computing and Networking*, pages 243–254, 2000.

[14] Evangelos Kranakis, Harvinder Singh, and Jorge Urrutia. Compass routing on geometric networks. In *CCCG*, 1999.

[15] Fabian Kuhn, Roger Wattenhofer, Yan Zhang, and Aaron Zollinger. Geometric ad-hoc routing: of theory and practice. In *PODC*, pages 63–72, 2003.

[16] Jinyang Li, John Jannotti, Douglas S. J. De Couto, David R. Karger, and Robert Morris. A scalable location service for geographic ad hoc routing. In *MOBICOM*, pages 120–130, 2000.

[17] Radhika Nagpal, Howard E. Shrobe, and Jonathan Bachrach. Organizing a global coordinate system from local information on an ad hoc sensor network. In *IPSN*, pages 333–348, 2003.

[18] D. Niculesu and B. Nath. Trajectory based forwarding and its applications. In *MOBICOM*, 2003.

[19] Goce Trajcevski, Oliviu Ghica, Peter Scheuermann, Marco Zuniga, Rene Schubotz, and Manfred Hauswirth. Improving the energy balance of field-based routing in wireless sensor networks. In *GLOBECOM*, 2010.

[20] L. F. W. van Hoesel and P.J.M. Havinga. Distributed coverage area reporting for wireless sensor networks, 2009. TR-CTIT-09-16.

[21] H. Wang, X. Zhang, and A. Khokhar. Efficient ”void” handling in contention-based geographic routing for wireless sensor networks. In *GLobecom*, 2007.

[22] X. Wu, G. Chen, and S.K. Das. Avoiding energy holes in wireless sensor networks with nonuniform node distribution. *IEEE Trans. Parallel and Distributed Computing*, 19(5), 2008.

[23] Jixing Yao, Guyu Zhang, Jinko Kanno, and Rastko Selmic. Decentralized detection and patching of coverage holes in wireless sensor networks. In *Proc. SPIE*, 2009.