ViBES: VIRTUAL BACKBONE FOR ENERGY SAVING IN WIRELESS SENSOR NETWORKS

Stefano Basagni Marco Elia Rituparna Ghosh
Department of Electrical and Computer Engineering
Northeastern University
E-mail: {basagni,melia,rghosh}@ece.neu.edu

ABSTRACT

Prolonged lifetime, robustness and scalability are important requirements in wireless sensor networks (WSNs). In this paper we show that via the energy-efficient construction of a connected backbone and a simple stateless routing over the backbone it is possible to achieve these desired goals. We present a ns2 simulation-based comparison between data dissemination in WSNs when a) each sensor routes the sensed data to the sinks directly (flat network organization), and b) a backbone is constructed on top of the flat network topology and only a few nodes are in charge of data delivery to the sinks (hierarchical organization, here termed ViBES: Virtual Backbone for Energy Saving). Multiple targets (i.e., moving nodes) are allowed to roam throughout the network of static sensors. Simulations are performed on networks of increasing densities where up to 3 targets are identified and reported to up to 3 sinks. Our results show that routing through the backbone increases the network lifetime up to seven times over the flat network organization where network lifetime is defined as a) the time till the first node dies because of energy depletion, b) the time required for a given percentage of the network nodes to die, and c) the time required for the network to get disconnected.

INTRODUCTION

Advances in MEMS and wireless communication technologies have recently paved the way and dramatically increased the opportunities for applications based on *wireless sensor networks* (WSNs). Monitoring possibly hostile environments as well as supporting peer to peer communications in the battlefield are now among the many applications that are possible via unmanned WSNs.

WSNs typically consist of large number of energy constrained, self-organizing wireless nodes forming multi-hop (ad hoc) networks. The nodes have sensing capabilities and the sensed data is routed to data collection points called *sinks* via intermediate sensors.

Figure 1 shows a typical sensor network, where nodes have been scattered throughout a geographic area and targets are roaming among them. A typical application is habitat monitoring, where every time a target enters the sensing area of a node, a corresponding packet is generated and routed to the closest sink.

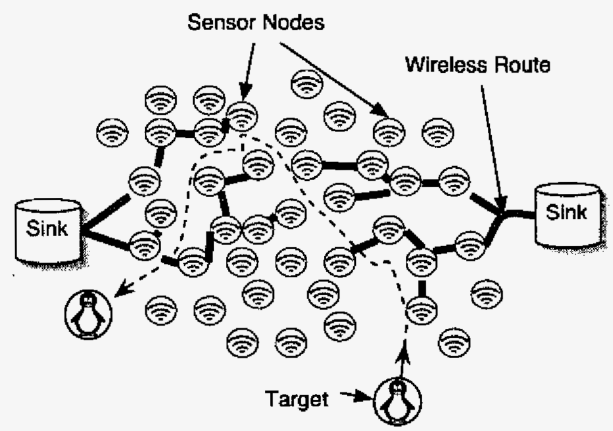

Figure 1: A WSN with sinks and targets.

It is widely recognized that one of the major barriers to the deployment of WSNs is the limited power at each node. Therefore, sensor network protocols should be designed with the aim of minimizing nodes' energy consumption and maximizing the overall network lifetime, i.e., the time till the network is able to provide the required service by delivering data to the sinks.

Typical definitions of network lifetime are the time from the start of network operation to (1) the time the first sensor node "dies," (2) the time a given percentage φ of the nodes "die," and (3) the time till the network is disconnected, i.e., there is at least a sensor node that has no route to any sink.

The first definition is common to most of previous research on WSN lifetime and it is intended to give an idea about how much, in general, a sensor can stay operational given the specific scenario it is in. In this sense, a network is considered "alive" until each of its sensors are. Definition (2)

0-7803-8847-X/04/$20.00 ©2004 IEEE

accounts for those applications in which the data collected are useful if received from a certain percentage of the nodes. The third definition is possibly the more natural one, given that the function of a sensor network is to deliver the sensed data to the sinks, which is possible only until the network is connected, i.e., until there is at least a route from each sensor to one of the sinks.

A host of protocols have been proposed for data dissemination is WSNs. Among these, directed diffusion [1], defines routing as *data centric* and clearly outlines the challenges and barriers of data dissemination protocol design. The effectiveness of directed diffusion has been thoroughly investigated showing a critical overhead of deploying this data dissemination scheme in a non-hierarchically organized network. In [2] enhancements to directed diffusion have been proposed where it is shown that, when it is deployed over a backbone, nodes' energy consumption is reduced to one third with respect to when the scheme is deployed over a flat network topology.

Specific protocols have been recently proposed that tackle explicitly with the issue of reducing energy consumption at each node and hence increase the overall lifetime of the network. Protocols like GAF [3] and ASCENT [4] partition the nodes into equivalent classes. Only a few nodes act as class representatives and provide routing capabilities at a certain time, while the other nodes are *turned off* to save energy. Network lifetime in GAF is defined as the time till the data delivery rate falls below a certain threshold. (Data delivery is defined as the ratio of the number of packets received to the number of packets sent.) Performance comparison with data delivery as obtained by using AODV as a routing protocol show up to a six-fold improvement on network lifetime. GAF needs location information, which is not always viable in WSNs. Network lifetime (definition (2) above with $\varphi = 90\%$) improvements in ASCENT are computed based on nodes being in different "ASCENT states" rather than on actual energy consumption. In this case, with respect to all nodes being constantly on, the improvement is around 300%.

A different approach has been adopted in [5] where a deterministic traffic scheduling algorithm balances the load over multiple paths between the sensor and the sink, which is proportional to the paths' residual energy. The authors report that, defining network lifetime as (1) above, the overall network lifetime increases of about 50% when their multipath routing is used instead of minimum energy routing.

Yet another approach is proposed in [6], where packet stream aggregation and traffic flow shaping are used for achieving longer network lifetime. Network lifetime is in accord to definition (1) above. The authors report that the lifetime of the network is increased by 90% with respect to the case with no stream aggregation.

Energy efficient clustering is used for energy saving by Heinzelman et al. in [7]. In LEACH the network is partitioned into clusters with one clusterhead and some ordinary nodes. A node becomes a clusterhead with a certain probability and surrounding nodes affiliate to it as ordinary nodes. Communication with the sink happens through the clusterhead that must be able to transmit directly to the sink (single-hop communication). Network lifetime is defined based on the time of death of the first node (definition (1) above) and on the time of death of the last node (definition (2), with $\varphi = 100$). With respect to the case in which all nodes (i.e., not just the clusterheads) are able to transmit their packets directly to the sink, LEACH shows improvements that are from 8 down to 3 times better. The major drawback of LEACH is that it assumes direct communication of any node with the sink (single-hop topologies).

A multi-hop clustering-based approach combined with an awake/asleep schedule has been recently proposed in [8]. An energy-conscious enhancement of the DMAC protocol proposed in [9], termed Sensor-DMAC (S-DMAC), is presented where a backbone among the network nodes is constructed in an energy efficient way. Nodes in the backbone are scheduled to stay awake, while all other nodes are instead sent to sleep for a predefined amount of time. Preliminary simulation results show that, with respect to DMAC, S-DMAC imposes a per-node power consumption which is from 40% (sparse networks) to 60% (dense networks) lower. (The measures are taken until the first node dies, i.e., according to definition (1).)

Our approach to energy saving and increased network lifetime makes use of the combination of an energy-effective *virtual backbone* construction together with a robust, simple, stateless routing scheme for data delivery. The idea is that only a subset of the sensor nodes that form a connected backbone is in charge of routing the data to the sinks, while all the other nodes have the sole task of sensing and sending the data to the closest backbone node. The sensed data eventually reaches the sinks via the backbone. A small fraction of the nodes are selected to be part of the backbone, and the actual backbone is then created by connecting the selected nodes via intermediate nodes and links. This "logical" connections being implemented via physical links among intermediate nodes, and the consistent energy savings obtained by routing over the backbone, motivated us to call the backbone itself *ViBES: Virtual Backbone for Energy Saving*.

Differently from previous solutions, the protocols we present here for ViBES construction and routing over the ViBES do not require extra hardware, or localization algorithm, to determine the nodes' location, and do not need each node to be a neighbor of the sinks.

We have demonstrated the effectiveness of using ViBES for energy-efficient data delivery to the sinks via thorough ns2-based simulations. We have simulated ViBES construction and routing in networks with up to 200 sensor nodes, where a maximum of three sinks are statically placed among the sensor, and up to three targets roam through the network and are monitored by the sensors. In particular, we have shown that with respect to robust, simple routing over the *flat* network (where *every* sensor participates to the data forwarding to the sinks), routing over ViBES prolongs network lifetime up to over 700%.

The paper is comprised of three main sections. The following section presents the protocols for ViBES construction and robust data delivery (routing to the sink) for both ViBES and the flat network. The third section describes the simulation environment and the simulation results. Conclusions and future research are outlined in the final section.

ViBES IN WSNs

In this section we describe the protocol for the construction of a ViBES and the routing used both on the flat networks and on the ViBES to deliver sensed data to the sinks.

ViBES Construction Protocol

Ideally we would want the sensor networks to live as long as possible. Traditional routing schemes where the sensed data is routed to the sinks over the flat topology (e.g., via shortest paths or other data centric techniques such as directed diffusion) are not always the best solution, because of their limited scalability and dependence on repeated floodings.

The idea behind ViBES is to use only a small subset of the nodes for data delivery to the sinks, saving the bulk of the nodes for just sensing the targets. The network is thus partitioned into the ViBES, composed of ViBES nodes, and ordinary nodes.

ViBES construction is comprised of two major phases: *primary ViBES node* selection and their interconnection to form a connected backbone. The selection of the ViBES nodes is performed at each node according to the Distributed Clustering Algorithm (DCA) proposed in [9]. Every node has a unique ID, a generic weight (which reflects the node's suitability for being a ViBES node, expressed, for instance, by its current amount of residual energy), and knows about the ID and the weight of all its one hop neighbors. Those nodes that have the biggest weight among their neighbors become primary ViBES nodes. The other nodes decide to be primary ViBES nodes or ordinary nodes based on the decision of all the neighbors with a bigger weight. Eventually, the process terminates with all the sensor nodes being partitioned into primary ViBES nodes and ordinary nodes.

A backbone is then constructed by connecting the primary ViBES nodes via some ordinary nodes, according to a simple rule first introduced in [10]. More precisely, primary ViBES nodes that are two or three hops apart, select common interconnection (ordinary) nodes which will be part of the backbone. The backbone paths thus formed guarantee that the obtained backbone is connected [10].

Figure 2 illustrates the process of selection of ViBES nodes. Every node in the flat network (top), based on local information about its immediate neighbors, decides whether to be a primary ViBES node or not (these nodes are the black squares in the bottom part of the figure).

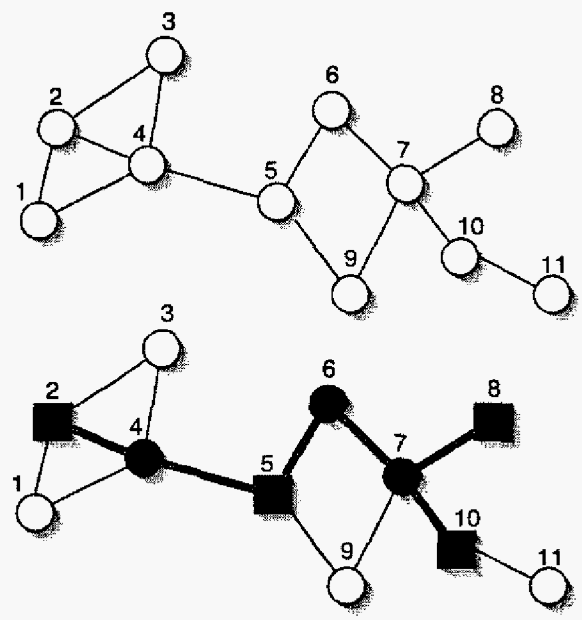

Figure 2: A flat network and its ViBES.

The final ViBES (bottom) is obtained by electing to the ViBES those nodes that interconnect the primary nodes. Details of how the primary nodes are chosen and then interconnected can be found in [11]. Notice that every sensor that does not belong to the ViBES has a ViBES node as neighbor.

Routing to the Sinks

The delivery of the sensed data to the sinks is performed in the following way. Once the network is up (i.e., after the initial neighbor discovery) each sink starts flooding out a control packet that carries its own ID and a counter that records the number of hops (distance) traveled by the packets. Upon receiving a packet, a node stores its distance from the sink that originated that packet, and forwards the packet after increasing the hop counter by one. This process also ensures that a node which is distant k hops from a sink knows who are all its neighbors at distance $k-1$. When a sensor has a data packet to transmit, it chooses the closest sinks and sends the packet to *all* its neighbors on the way to that sink. (In case of equidistant sinks, one is chosen randomly.)

While all the nodes in the flat network receive the flooding packets from the sink, when using a ViBES the flooding is restricted only to the ViBES nodes. Thus, routing on the flat network involves all nodes, while in the ViBES case, a sensor that is not on the ViBES sends its data to one of its neighboring ViBES nodes (i.e., routing happens only through the ViBES).

This simple, stateless routing has been chosen mainly for the following two reasons, which fit the requirements of protocol design for sensor networks.

- It is robust, in the sense that it provides multiple routes from each node to each sink, even on the ViBES. When one of the routes fails because of the failure of an intermediate node, the packet is still traveling through alternate paths.
- Beside the initial, "una tantum" flooding, this routing has no overhead. There is no route management, route discovery or routing table exchange for finding and maintaining a route to a sink.

SIMULATION RESULTS

We have demonstrated the effectiveness of data delivery to the sinks through the backbone by simulations. The ViBES construction protocol, and the simple routing both over the ViBES and over the flat network have been implemented in the VINT project network simulator "ns2" [12]. All phases of the described protocols have been implemented, which include the initial neighbor discovery (we assume all network nodes are turned on during a time interval $\tau = 1$s), routing table establishment and backbone construction.

Ns2 Settings and Experiments

Our implementation is based on the CMU wireless extension to ns2, i.e., we use a IEEE 802.11-like MAC with the DCF for our investigation. The sole parameter we have modified is a node maximum transmission range (sensor nodes are not likely to have the full IEEE 802.11 range, mainly for power consumption reasons).

All simulations refer to static wireless scenarios, in which n wireless nodes, with maximum transmission radius of 30m, are randomly and uniformly scattered in a geographic area which is a square of side L.

Each device has initial (residual) energy equal to 2J. The power consumed while transmitting, receiving, and while in the idle state are 175mW, 175mW and 0.015mW, respectively. The power consumed for sensing is equal to 1.75μW, and the sensing range is 10m. These values are consistent with those described in the specifications of the μAMPS-1 sensor, developed as part of the μAMPS project at M.I.T. [13].

In the simulation results presented here, the number of nodes n has been assigned values from 50 to 200 with increments of 25. The area side L has been set to 200m. This allowed us to test the advantages of using ViBES on increasingly dense networks, from (moderately) sparse networks ($n = 50$), with average degree equal to 3.5 to dense networks ($n = 200$) where the average degree equals 12.3. The number of sinks and the number of target is variable. The experiments here refer to four different scenarios: (a) one sink and one target; (b) three sinks and one target; (c) one sink and three targets, and (d) three sinks and three targets. Sinks are static, and have no resource constraints. They have the same transmission radius of a sensor node. They are placed randomly and uniformly among the sensor nodes. The targets move according to the random waypoint mobility model at a velocity of 5m/s and a maximum pause time of 5s (small animals like squirrels and fast penguins have similar speeds). The target nodes beacon their presence every 5s, with a "communication distance" of 10m. This mimics the sensing: Whenever a sensor receives the beacon of a target, a packet is created and routed to the closer sink, as explained in the previous section. (The cost of this "sensing" is considerably less than radio transmission/reception.)

All results are obtained by averaging over 1000 topologies for each scenario. This guaranteed a confidence interval of 95% and a 5% precision most of the times.

Network Lifetime

The main metric we are interested in investigating in this paper is the network lifetime. In particular, as mentioned, data delivery to the sink both on the flat network and over the ViBES has been performed and the network lifetime has been measured for both cases in the four described scenarios.

We consider three different definitions of network lifetime. Network lifetime is the time from the start of network operations to (1) the time the first sensor node "dies" because of energy depletion, (2) the time the 30% of the nodes die, and (3) the time till the network is disconnected, i.e., there is at least a sensor node that has no path to any of the sinks. In this paper we have chosen the percentage φ of the dead nodes that define the network lifetime according to definition (2) to be 30 because we have observed that, on average, greater values disconnect the network (this case is covered by definition (3)).

(a) One sink and one target

Figure 3 shows the improvement on network lifetime when using ViBES as opposed to routing on the flat network in the case of one sink and one target roaming through the network. The fact that in general, the lifetime decreases with the increase of the number of nodes accounts for the increasing density and thus increased routing messages.

We observe that independently of the particular definition of network lifetime, routing on ViBES increases the lifetime up to 735% (lifetime definition (1)), 165% (definition (2)) and 335% (definition (3)). These improvements are all obtained for the denser scenarios ($n = 200$). The reason for which the least improvement is obtained for definition (2) has to do with the fact that in a very dense network, the 30% of the nodes (60 nodes) in the flat network can die and the network is still connected, while only the death of 20 nodes, on average, disconnects the ViBES (this "strongness" result has also been obtained via simulations). In sparser networks, as expected, the improvement, although always present, is more modest, since most of the network nodes are in the ViBES.

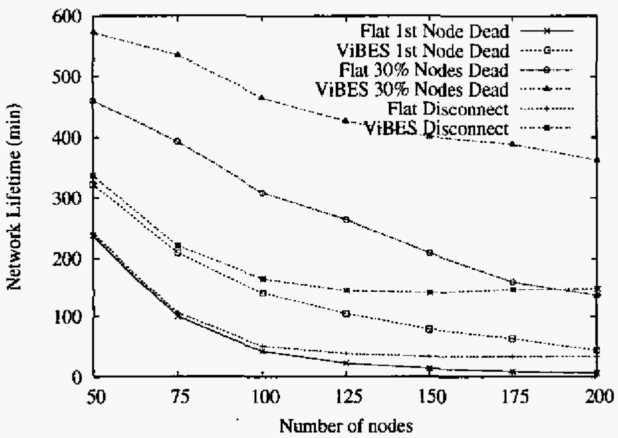

Figure 3: Network lifetime: One sink and one target.

(b) Three sinks and one target

Figure 4 shows the network lifetimes when there are three sinks randomly placed in the network and only one target roaming around. In this case the "peak improvement" is not as impressive as for the scenario above. Routing on ViBES induces lifetime improvements of 174%, 88% and 87% for the three definitions. Having multiple sinks in place, induces higher energy savings, since the average route length to the closest sink is now lower. This, however, given the higher number of routes in the flat topology, is more helpful for routing in the flat network than for routing on ViBES. The routes to the sinks over the backbone are generally longer than those on the flat network. The difference is not remarkable, and it is less evident in case of longer routes, like those of scenario (a). Having generally shorter routes now makes the difference between "flat routes" and "ViBES routes" more important, which accounts for the lower improvements.

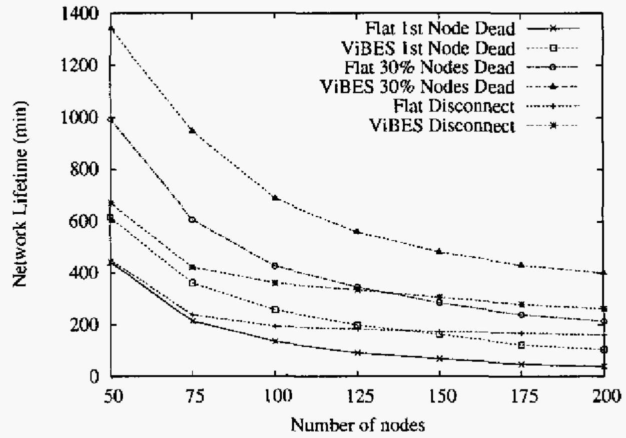

Figure 4: Network lifetime: Three sinks and one target.

(c) One sink and three targets

Figure 5 shows the scenario with one sink and three moving targets. All measures of the lifetime are noticeably reduced when there are three targets roaming. This is expected as the three targets are continuously creating traffic throughout the network and draining the energy of the nodes (the time the generic node spend in idle state is very short compared to when it is transmitting or receiving). In this scenario the best improvements of using ViBES are of 268%, 159% and 142% for the time the first node dies, the time till 30% of the nodes die, and the time the network becomes disconnected, respectively. Again, these improvements are significantly reduced compared to scenario (a). This is by far the scenario that offers the least improvements: Routes are generally longer (there is only one sink) and always busy (the three targets generate a consistently high amount of traffic).

(d) Three sinks and three targets

Results for the case with three targets roaming in a network with three sinks, are depicted in Figure 6. Compared to the same number of targets, but only one sink in place (case (c)), a longer lifetime is achieved, as expected, given that the shorter routes to the sinks now require less energy for data delivery. The ViBES gives peak improvements of 349%, 232% and 115% for the three different definitions of network lifetime.

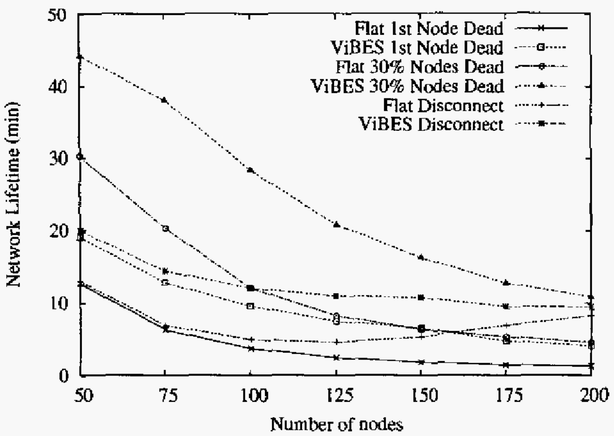

Figure 5: Network lifetime: One sink and three targets.

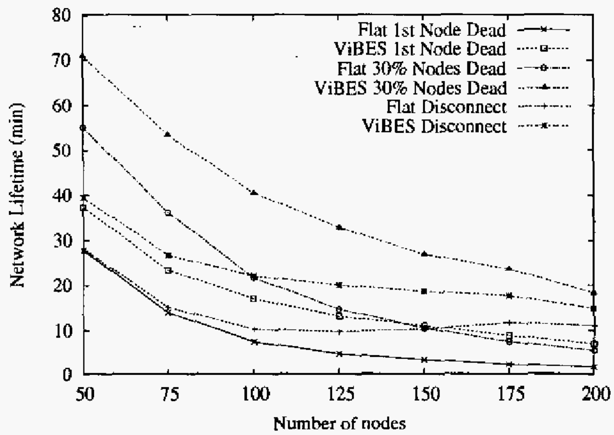

Figure 6: Network lifetime: Three sinks and three targets.

To conclude this section we notice that the use of ViBES is extremely effective in prolonging the network lifetime also because the energy consumed for ViBES construction is very modest. Simulation results have shown that the average energy consumption per node for setting up the ViBES is always less than 0.2J in all the considered scenarios. Thus, when the targets start roaming, the node energy is basically the same both in the flat and in the ViBES case.

CONCLUSIONS AND FUTURE RESEARCH

In this paper we have shown that routing data to the sinks of a WSN via ViBES helps in obtaining consistent energy savings and overall network lifetime improvement. Using three different lifetime definitions: the time required for the first node to die, the time required for 30% of the nodes to die, and the time required for the network to get disconnected, we have shown improvements on network lifetime of up to over 700%.

Simulations on networks of various densities and varying number of sinks and targets have allowed us to observe that the ViBES approach works best in low/medium traffic conditions. However, despite the extra communication needed for ViBES construction, routing over ViBES always outlasts routing over the flat topology in all cases.

With the present paper our main aim has been to demonstrate that a hierarchical organization of the sensor nodes can lead to significant energy savings for WSNs. Here, our metric of interest has been network lifetime, and the result obtained are definitely encouraging.

More needs to be done, in several different directions, which include 1) dealing with the mobility of the sensor nodes, or with the mobility of the sinks, 2) being able to make ViBES adaptive to the changing node status (e.g., keep on the ViBES the nodes with the higher residual energy), and to the addition of new, fresh nodes, 3) demonstrate the effectiveness of routing over ViBES with respect to other metrics, such as the data delivery ratio, and 4) optimize the positioning of sinks, evaluating how this affects network lifetime and data delivery ratio.

ACKNOWLEDGMENTS

This research has been partially performed while the first author was wisiting Professor Chiara Petrioli at the University of Rome "La Sapienza," supported by the project IST-2001-34734 EYES (EnergY Efficient Sensor networks).

References

[1] C. Intanagonwiwat, R. Govindan, D. Estrin, J. Heidemann, and F. Silva. Directed diffusion for wireless sensor networking. *IEEE/ACM Transactions on Networking*, 11(1):2–16, February 2003.

[2] A. Marcucci, M. Nati, C. Petrioli, and A. Vitaletti. Hierarchical routing and data aggregation in wireless sensor networks. Technical report, Università di Roma "La Sapienza," Roma, Italy, October 2003.

[3] Y. Hu, J. Heidemann, and D. Estrin. Geography-informed energy conservation for ad hoc routing. In *Proceedings of the 7th ACM Annual International Conference on Mobile Computing and Networking*, pages 70–84, Rome, Italy, July 16–21 2001.

[4] A. Cerpa and D. Estrin. ASCENT: Adaptive Self-Configuring sEnsor Networks Topologies. In *Proceedings of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies, INFOCOM*

2002, volume 3, pages 1278–1287, New York, NY, June 23–27 2002.

[5] N. Jain, D. K. Madathil, and D. P. Agrawal. Energy aware multi-path routing for uniform resource utilization in sensor networks. In *Proc. of the 2nd International Workshop on Information Processing in Sensor Networks, IPSN 2003*, volume 2634 of *Springer LNCS*, pages 473–487, Palo Alto, CA, April 22–23 2003.

[6] C. Schurgers and M. B. Srivastava. Energy efficient routing in wireless sensor networks. In *Proceedings of IEEE MILCOM 2001*, pages 357–361, Vienna, VA, October 28–31 2001.

[7] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan. Energy-efficient communication protocol for wireless microsensor networks. In *Proceedings of the 33rd Annual Hawaii International Conference on System Sciences, HICSS 2000*, pages 3005–3014, Maui, HA, January 4–7 2000.

[8] S. Basagni, A. Carosi, and C. Petrioli. Sensor-DMAC: Dynamic topology control for wireless sensor network. In *Proc. of IEEE VTC 2004 Fall*, Los Angles, CA, September 26–29 2004.

[9] S. Basagni. Distributed clustering for ad hoc networks. In A. Y. Zomaya, D. F. Hsu, O. Ibarra, S. Origuchi, D. Nassimi, and M. Palis, editors, *Proceedings of the 1999 International Symposium on Parallel Architectures, Algorithms, and Networks (I-SPAN'99)*, pages 310–315, Perth/Fremantle, Australia, June 23–25 1999. IEEE Computer Society.

[10] I. Chlamtac and A. Faragó. A new approach to the design and analysis of peer-to-peer mobile networks. *Wireless Networks*, 5(3):149–156, May 1999.

[11] S. Basagni, D. Turgut, and S. K. Das. Mobility-adaptive protocols for managing large ad hoc networks. In *Proceedings of the IEEE International Conference on Communications, ICC 2001*, volume 5, pages 1539–1543, Helsinki, Finland, June 11–14 2001.

[12] The VINT Project. *The ns Manual*. http://www.isi.edu/nsnam/ns/, 2002.

[13] R. Min and A. Chandrakasan. Energy-efficient communication for ad-hoc wireless sensor networks. In *35th Asilomar Conference on Signals, Systems, and Computers*, volume 1, pages 139–143, Pacific Grove, CA, November 4–7 2001.