

A Singlehop Collaborative Feedback Primitive for Wireless Sensor Networks

Murat Demirbas

Onur Soysal

Muzammil Hussain

Department of Computer Science and Engineering
University at Buffalo, SUNY, Buffalo, NY, 14260
{ demirbas | osoysal | mh69 }@buffalo.edu

Abstract—To achieve scalability, energy-efficiency, and timeliness, wireless sensor network deployments increasingly employ in-network processing. In this paper, we identify *singlehop feedback collection* as a key building block for in-network processing applications, and introduce a basic singlehop primitive, *pollcast*. The key idea behind this primitive is to exploit the *receiver-side collision detection* information at the MAC-layer to speed-up collaborative feedback collection. Using *pollcast*, a node can get an affirmation about the existence of a node-level predicate P in its neighborhood in constant time by asking all nodes where P hold to reply simultaneously. We have implemented *pollcast* on Tmotes using Chipcon 2420 radio. Our results show that this primitive is indeed lightweight, resilient, and effective. Our paper is also the first time receiver-side collision detection is achieved in a practical manner for Chipcon 2420 radio.

I. INTRODUCTION

Early deployments of wireless sensor networks (WSNs) have been mostly limited to passive data collection, where sensor readings from the network are relayed toward a basestation for storage and processing [1], [2]. In order to cope with the bandwidth/energy consumption and latency concerns associated with this centralized approach, *in-network information processing* has been widely adopted [3]–[5]. In-network processing exploits the computation capability of sensor nodes to process data locally in the network close to where it originates. Two main use-cases of this approach are 1) to summarize and reduce the data transmitted to the basestation, and 2) to perform decisions locally to avoid contacting the basestation for each decision. An example of the first is the aggregation of data and filtering of false-positives and duplicates, reducing the energy wasted in routing these packets all the way to the basestation only to be disposed there. An example of the second appears in intruder-interceptor applications, where having a leader responsible for the current intruder detection is critical for maintaining a tracking structure [6]–[8].¹

Although there have been many efficient “point solutions” [3], [4], [10] to the problems that appear in the context of these two use-cases, there has been no effort to address the issue by developing efficient and general collaboration primitives. In contrast to devising point solutions, designing

¹In this scenario, in-network processing is required even for correctness, as it has been shown that for satisfying optimality constraints, the latency with which an interceptor requires information about the intruder it is tracking depends on the relative locations of the two: the closer the distance, the smaller the latency [9].

such primitives would improve reusability and integration, and provide a unified framework for standardization of in-network processing protocols.

Contributions and overview. We identify collaborative feedback collection from a singlehop neighborhood as a key building block for both use-cases of in-network processing. We observe that wireless broadcast has many useful features for facilitating collaboration. Firstly, broadcasting is atomic: that is, all recipients receive a broadcast at the same time, which is useful for synchronizing the nodes in singlehop for building a structured feedback collection. Secondly, broadcast allows receiver-side collision detection (RCD): that is, a snooping node can detect collisions of messages, which is useful for extracting feedback from multiple nodes in a quick manner. Exploiting these features of wireless broadcast, we propose an efficient and lightweight singlehop collaborative feedback primitive, *pollcast*.

Using *pollcast*, a node can take a quick poll from its neighborhood by asking all nodes with a certain property to reply simultaneously. An operation starts with a “poll” phase, where the initiator broadcasts a poll message of the form “does there exist any node with property P ?”. The initiator then moves on to the second phase, the “vote” phase, to listen for the responses to its poll. If the initiator hears silence in the vote phase, it concludes that the polled predicate P does not hold for any node. Otherwise, if there is a response or there are multiple responses (in which case the initiator detects a collision via RCD), the initiator concludes that P holds for some nodes. Thus, regardless of the number of nodes that need to reply, *pollcast* completes in $O(1)$ time. Some applications of *pollcast* operation are false-positive suppression, clustering, and the querying of the neighborhood for debugging purposes.

We have implemented *pollcast* on the Tmotes [11] using the popular Chipcon 2420 radios. We show performance results from this implementation. In addition, we have implemented *pollcast* under a WSN simulator [12] to be able to perform more controlled experiments for a larger scale networks and compare our primitive with other protocols.

Finally, our paper is the first time RCD feedback is achieved for the popular CC2420 radios in a practical manner. Our experiments indicate that our RCD implementation has around 100% completeness and 0% false-positive detections. Our RCD technique is easily achievable at the MAC layer in

software and does not require any modification to the physical layer or the wireless radio hardware.

Singlehop wireless broadcast has recently been identified as a narrow-waist suitable for standardization efforts in the WSNs [13]. Our efficient and lightweight singlehop collaborative feedback primitives for supporting in-network processing will help boost these standardization efforts. Researchers working on cooperative control may become end-users of our primitives, since pollcast is suitable for a control-theoretic framework of periodically collecting information about the state of the system and imposing a compensated control over the system.

II. RELATED WORK

Collision detection. The feasibility of collision detection for CC1000 (mica2) radios has been demonstrated in [14] for a limited context (for certain capture/shadowing effect scenarios). The success rate of the preamble-based collision detection used in [14] drops quickly for more than two simultaneous senders. Our RCD implementation is based on receiver-side carrier sensing and is more general and inclusive than preamble-based RCD. Several existing MAC layers, such as B-MAC [15], already support the carrier-sensing capability required for our collision detector.

A recent empirical study on CC1000 radios [16] linked the successful message reception in the presence of concurrent transmissions to the signal-to-interference-plus-noise-ratio (SINR) exceeding a critical threshold. The results also showed that it becomes harder to estimate the level of interference in the presence of multiple interferers, and that the measured SINR threshold generally increases with the number of interferers. Hence, CC1000 performs poorly in the presence of more than a couple of concurrent transmitters. Our preliminary experiments find that the CC2420 radio behavior under concurrent transmissions is more resilient. CC2420 radio is able to receive a message successfully among a set of several concurrently transmitted messages, due to the direct-sequence spread spectrum (DSSS). Thus, in CC2420 radio, SINR threshold for successful message reception becomes harder to define and it becomes harder to correlate RSSI with the successful message reception.

Singlehop programming abstractions. Several programming abstractions have been proposed for WSNs [17]–[20].

The Tenet project [17] proposes a tiered WSN architecture with small-form-factor motes and more powerful master nodes. Tenet asserts that complex application logic should be implemented only on the masters. Applications running on masters task motes, and motes just communicate back to the masters the results from these tasks. Our singlehop polling primitive would be instrumental for the masters in Tenet to achieve quick and ad hoc feedback collection from the motes ondemand.

In an effort to simplify the adoption of distributed algorithms for WSNs in terms of a neighborhood abstraction, Hood [18] provides an API that facilitates exchanging information among a node and its neighbors. For example,

Hood can define a one-hop neighborhood over which light readings are shared. Beneath the API, Hood automatically discovers neighbors and caches the values of their attributes periodically, while simultaneously sharing the values of the node’s own attributes. Similar to Hood, abstract regions [19] and TeenyLime [20] propose mechanism for discovery and sharing of data (structured as tuples) among sensor nodes.

Using the information exchange mechanisms proposed in these abstractions [18]–[20], it may be possible to achieve a constant response time to a query by performing periodic state exchange among neighbors behind the curtains. However, a big problem facing these approaches is to decide on the frequency of this exchange. If the exchange is done infrequently, the query will be answered using stale data. (This is especially problematic for real-time applications such as intruder-interceptor applications.) If the exchange is done frequently, a lot of traffic is generated wasting precious energy and bandwidth. In contrast to these work that deal with state exchange among nodes, our focus in pollcast is to provide a lightweight and efficient framework for on-demand binary feedback collection from neighbors.

III. RECEIVER-SIDE COLLISION DETECTION

A. RCD Implementation

Below we discuss pros and cons of three possible approaches to RCD implementation.

Received-signal-strength-indicator (RSSI) based collision detection depends on monitoring the RSSI information from the radio frequently (i.e., for every byte) and looking for patterns that imply the existence of a collision. RSSI-based RCD is a low-level and general technique for collision detection, but it brings additional processing burden (due to the frequent interrupts it generates for RSSI processing). Moreover, compared to the CC1000 radios where the correlation with interference and SINR is observable [16], achieving RSSI based collision detection is much harder for the CC2420 radios.

Cyclic-redundancy-check (CRC) based collision detection depends on checking the CRC bits of received messages, and raising a collision detection upon encountering a bad CRC bit. CRC is reliable, however, CRC-based RCD is applicable only when the radio is locked to a certain message and preamble and packet frame are received. Thus, CRC based RCD is not general enough for detecting all type of collisions.

Carrier sensing based collision detection depends on sensing the medium for ongoing transmissions. While carrier sensing is widely used by transmitters in wireless networks with CSMA MAC layers (including IEEE 802.11, IEEE 802.15.4, and all of the WSN MAC protocols), we adopt this technique for use by the receiver to detect collisions. The CC2420 radio exports a Clear-Channel-Assessment (CCA) signal for the purposes of carrier-sensing. CCA is well-engineered and robust, and is calculated by the radio chip based on a window of RSSI readings and thresholding. Since CCA has the additional benefit of radio-level support, CCA obviates the need to involve the CPU in collision detection, and is simpler and more




Fig. 1. Collision detection performance

reliable than RSSI. The challenge with CCA signal is that, it is generated only after a transmission is scheduled to the radio, so for implementing RCD using CCA we manipulate the CC2420 radio to perform CCA also in the idle state. Using CCA-based collision detection we are able to detect collisions even when no intelligible packet information (such as the preambles) is received due to interference.

We implement both CRC- and CCA-based RCD for CC2420 radios and compare their performances next.

B. RCD Experiments

To evaluate our RCD implementation, we use a setup consisting of upto 6 Tmote-Invent nodes with CC2420 radios. One of these nodes is designated as the poller, and the remaining nodes are programmed as voters. In each experiment, the poller broadcasts a start message, upon which the voters transmit a reply immediately (without performing CSMA as per our modification to the TinyOS MAC layer). We initially start with two voters and repeat each experiment 200 times, before increasing the number of voters by one to investigate the RCD performance under a more contended scenario. Figure 1 presents our findings.

CRC-based RCD combined with successful reception of one of the votes achieves upto 100% detections in the 2 and 3 voters case. In the 4 and 5 voter cases, the shadowing effect decreases as the power-sum of other transmissions create excessive interference for successful reception of any single message. So CRC-based detections and successful receptions of one vote decrease for these cases, since the probability of successfully locking to a preamble decreases with the degradation in shadowing effects. These results show that CRC-based RCD is insufficient for detecting most collisions.

CCA-based RCD achieves a thorough detection of collisions and is not noticeably affected by multiple voters as Figure 1 shows. We repeat CCA-based RCD experiments with varying the distance between the nodes and achieve similar very good completeness results for collision detection.

Finally, we test the false-positive ratio of our RCD implementation by repeating the experiments with 0 voters. In our experiments, CCA-based detection has reported 1 false positive out of 200 tries, so we conclude that the false-positives do not constitute a problem for our RCD implementation.

IV. POLLCAST OPERATION

A. Pollcast Implementation

As we outlined in the Introduction, our pollcast operation consists of two consecutive phases: poll-vote. In the poll phase, an initiator may start a pollcast operation for predicate P by broadcasting a $POLL_P$ message. In the vote phase, the initiator switches to listen for the votes: receiving a VOTE or a collision implies that there exist some nodes with the polled property P , whereas silence means that no node satisfies P . Any node receiving a $POLL_P$ message in the poll phase should vote accordingly in the vote phase. A VOTE broadcast is performed only if the answer is “yes” for the node-level predicate P . The nodes communicate a “no” by keeping quiet in the vote phase.

The poller can limit its poll to a subset of the neighborhood by including a list of intended participants in the $POLL_P$ message (otherwise, all neighbors are included in the polling by default). This participant list feature is useful because, using the result of a previous poll, the poller may adaptively select a subset of participants for its next poll to narrow in on a clue. This way incremental searching for a condition may be possible, but we do not explore this any further here.

Pollcast operations in multihop networks. Even though pollcasts are singlehop operations, when they are executed in a multihop network, interference from neighboring regions in the form of hidden terminal problems are unavoidable. For the sake of simplicity, we assume “atomicity” of pollcasts in a 2-hop region. More specifically, we assume that when a pollcast is in progress, there cannot be another simultaneous pollcast within 2-hops of this pollcast. This assumption is motivated by practical reasons. Using this assumption we are able to keep our pollcast implementation extremely lightweight and short in duration. In contrast, trying to cope with collision of poll messages would require either a globally synchronized rounds approach (as in [21], [22]) or introduction of several new control messages (such as an RTS/CTS handshake with every neighbor as in [23]) and the overhead may defeat the purpose: lightweight singlehop collaborative feedback!

Our atomicity assumption is reasonable for low traffic WSN deployments. In most WSN deployments the network is idle for most of the time. For the case of bursty triggering of the pollcasts (say, for example, due to detection of an intruder in the area), we rely on CSMA to arbitrate our lightweight pollcast operations in singlehop. Although we assume atomicity in the design and presentation of pollcast, our experiment and simulation results investigate the effects of collisions and hidden terminal problems on the consistency of pollcasts under low, medium, and heavy traffic loads.

B. Pollcast Applications

Pollcast enables a node to corroborate its estimates about the environment with the neighboring nodes. One practical example is in suppressing false-positive detections. Due to inexpensive sensors, false-positives are frequent occurrences in WSN deployments. For example, in the “Line in the sand” [2] and “ExScal” [24] applications, sensors would often have

false-positive detections (due to heat drifts and sunlight in out-door environments for PIR and due to noise for magnetometer). False-positives cause problems as they use up the precious resources on the network. A false positive detection is forwarded over several hops until it reaches a basestation node or a clusterhead which can drop this detection as an outlier. Pollcast operation is useful for filtering out false-positive detections effectively. A node that has a detection (and that has not yet heard any other detection from neighboring nodes) may incite a pollcast operation to ask neighbors if they also detected the alleged phenomena P . This is a yes/no question and can be executed very quickly using pollcast. Only if the answer is affirmative (i.e., there is activity in the vote phase) the initiator of the poll notifies the basestation about its detection.

Other applications of pollcast follow from the second use-case scenario of in-network processing, namely in-network control. A good example for this case is barrier synchronization. In WSNs barrier synchronization is useful for synchronizing operations in each cluster. For example, the clusterhead can query the member nodes as to whether they are finished with the current phase of the operation before the cluster as a whole can move to the next phase in the operation. Another example of in-network control is leader election. Even though an object is detected by many nodes at the same time, it is important to elect a leader primarily responsible for the object. The leader can then process and send messages, keep track of the trajectory of the object, and hand it off to the next leader. It may be advantageous to select the leader based on the strength of its detection. This can be checked and established simultaneously, by making the initiator do a pollcast by announcing its detection and challenging anyone with a greater detection to vote. If no node votes, the initiator announces its leadership at the end. Electing leaders is also important in clustering applications. Several metrics may be used in the election of the clusterhead. For example, an initiator may check to ensure that there are no other clusters (clustermembers) in its singlehop [10].

C. Pollcast Experiments

For performing large-scale and controlled experiments, we have implemented pollcast in Prowler [12], a MATLAB based event driven simulator for distributed WSNs. We use Prowler to simulate the radio transmission/propagation/reception delays of Tmotes, including collisions in WSNs.

We compare pollcast with a strawman polling protocol *naivepoll*. In *naivepoll*, after the initiator broadcasts a poll message for a predicate P , the neighbors where P hold send their votes via CSMA. This, of course, means that if all the voters are not within singlehop of each other, there may be collisions of votes at the poller. Unlike pollcast, *naivepoll* does not employ RCD to learn from collision of votes.

We simulate pollcast over a 10x10 grid of 100 nodes. Each node has 8 neighbors (except the boundary nodes which may have 3-5 neighbors). We perform the simulations varying the number of concurrent pollers from 1 to 10 in unit steps and

then incrementing from 20 to 60 in steps of 10. All elected pollers for a given simulation run try to execute their polls in the beginning of the run, with only CSMA arbitrating between these pollers. For each simulation run, we also randomly initialize the P predicates at the nodes, so that the number of voters for a poll vary randomly.

Settling time of a simulation run is defined as the duration between the first and last message sent in the simulation run. The settling time graph in Figure 2 show that pollcast takes less time to complete than *naivepoll*. This is expected as pollcast uses the backoffs in CSMA only for the poll messages, but in *naivepoll* there are backoffs for both poll and vote messages. A pollcast operation in isolation is about 3msecs, and as the figure shows, for low traffic case the settling time of pollcasts in the WSN is around 3 to 5 msecs.

Figure 3 investigates the loss of poll messages at the receiver nodes and finds that a significant number of poll messages are lost in both protocols due to the hidden terminal problem. Pollcast seems to lose more poll messages because it sends vote messages without any carrier sensing. So in pollcast a vote message may be sent even when another poll message is being transmitted in the neighborhood.

In Figure 4 we compare *naivepoll* and pollcast for inconsistencies: when no vote is received at the poller whereas in reality the poller has at least one neighbor where P holds. In *naivepoll* an inconsistency occurs when all votes collide at the poller. Since *naivepoll* does not employ RCD, it cannot deduce the *existence of at least one vote* in this case. In pollcast an inconsistency occurs only when the poll message is lost on *ALL* the voters. When comparing Figure 4 with Figure 3 for the pollcast protocol, one should keep in mind that the loss of poll messages in Figure 3 is for *any* neighbor, while Figure 4 indicates the loss of poll messages on *all* the voters. Hence we find that the inconsistencies for pollcast is always less than 1/3 of that of *naivepoll*. Especially for the low traffic case, the inconsistencies in pollcast is almost negligible.

Figure 4 also shows the number of false positives for pollcast. False positives occur when a poller detects at least one vote (collision of votes) when in fact there were no votes for it. This may happen in pollcast because in the high traffic case a poller may misinterpret a collision of two other polls as a vote for its poll. Our experiments show that the number of false-positives for pollcast is very low.

V. CONCLUDING REMARKS

In this paper, we presented a lightweight and efficient singlehop primitives for collaborative feedback collection. Using pollcast a node can query its singlehop neighborhood about a predicate P and in $O(1)$ time learns whether there are neighbors for which P holds. Our proposed RCD technique for the implementation of pollcast operation is easily achievable at the MAC layer in software and does not require any modification to the physical layer or the wireless radio hardware. As such, the pollcast operation is readily applicable in WSNs and in a more general context for building robust mobile ad hoc network applications over 802.11-

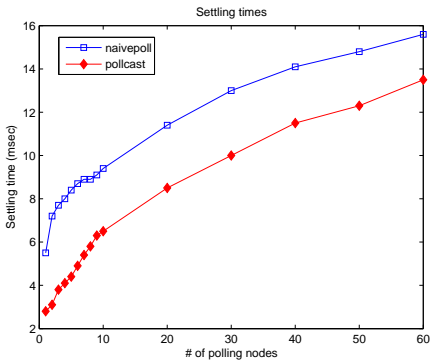


Fig. 2. Settling time

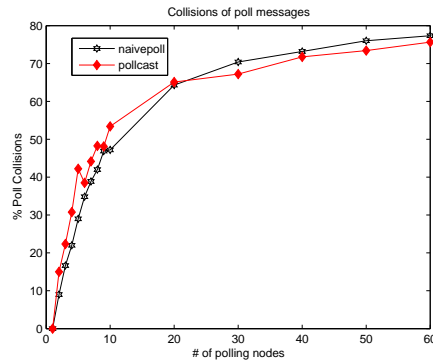


Fig. 3. Poll message collisions




Fig. 4. Polling inconsistencies

enabled networks. Some applications of pollcast in WSNs are in efficient implementations of false-positive filtering, in-network intruder classification, clustering, leader election, and barrier synchronization algorithms. A desirable and obvious enhancement that we added to the pollcast operation is to return an approximate count of the number C_P of nodes a property P holds. This enhancement would be useful for querying of the neighborhood for classification of an intruder (say as a soldier, car, or tank as in Exscal [24]) by counting the detections in the neighborhood. Due to limited space we refer reader to the technical report version of this study [25] for this discussion.

Another interesting topic for future research is to determine which lowerbounds would apply for common tasks (such as majority detection, at-least- k detection, etc.) developed in our framework. Since the poller can use the results from a previous poll to adaptively select a subset of its neighbors for its next poll, it is possible to devise several efficient strategies for these tasks.

REFERENCES

- [1] G. Tolle, J. Polastre, R. Szewczyk, N. Turner, K. Tu, P. Buonadonna, S. Burgess, D. Gay, W. Hong, T. Dawson, and D. Culler, "A macroscope in the redwoods," in *SenSys*, 2005.
- [2] A. Arora and et. al., "A line in the sand: A wireless sensor network for target detection, classification, and tracking," *Computer Networks (Elsevier)*, vol. 46, no. 5, pp. 605–634, 2004.
- [3] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "A survey on sensor networks," *IEEE Communications Magazine*, 2002.
- [4] C. Intanagonwiwat, R. Govindan, D. Estrin, J. Heidemann, and F. Silva, "Directed diffusion for wireless sensor networking," *IEEE/ACM Trans. Netw.*, vol. 11, no. 1, pp. 2–16, 2003.
- [5] D. Estrin, R. Govindan, J. S. Heidemann, and S. Kumar, "Next century challenges: Scalable coordination in sensor networks," in *Mobile Computing and Networking*, 1999, pp. 263–270.
- [6] M. Demirbas, A. Arora, and M. Gouda, "Pursuer-evader tracking in sensor networks," *Sensor Network Operations*, IEEE Press, 2006.
- [7] M. Demirbas, A. Arora, T. Nolte, and N. Lynch, "A hierarchy-based fault-local stabilizing algorithm for tracking in sensor networks," *OPDIS*, pp. 299–315, 2004.
- [8] V. Kulathumani, M. Demirbas, A. Arora, and M. Sridharan, "Trail: A distance sensitive wireless sensor network service for distributed object tracking," *EWSN*, pp. 83–100, 2007.
- [9] H. Cao, E. Ertin, V. Kulathumani, M. Sridharan, and A. Arora, "Differential games in large scale sensor actuator networks," in *IPSN*, 2006, pp. 77–84.
- [10] M. Demirbas, A. Arora, V. Mittal, and V. Kulathumani, "A fault-local self-stabilizing clustering service for wireless ad hoc networks," *IEEE Trans. on Parallel and Distributed Systems*, vol. 17, no. 9, pp. 912–923, 2006.
- [11] "Moteiv," <http://www.moteiv.com/>.
- [12] G. Simon, P. Volgyesi, M. Maroti, and A. Ledeczi, "Simulation-based optimization of communication protocols for large-scale wireless sensor networks," *IEEE Aerospace Conference*, pp. 255–267, March 2003.
- [13] D. Culler, P. Dutta, C. T. Ee, R. Fonseca, J. Hui, P. Levis, J. Polastre, S. Shenker, I. Stoica, G. Tolle, and J. Zhao, "Towards a sensor network architecture: Lowering the waistline," *HotOS X*, 2005.
- [14] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, and D. Culler, "Exploiting the capture effect for collision detection and recovery," *EmNetS*, May 2005.
- [15] J. Polastre, J. Hill, and D. Culler, "Versatile low power media access for wireless sensor networks," in *SenSys*, 2004, pp. 95–107.
- [16] D. Son, B. Krishnamachari, and J. Heidemann, "Experimental study of concurrent transmission in wireless sensor networks," in *Conference on Embedded networked sensor systems*, 2006, pp. 237–250.
- [17] O. Gnawali, K.-Y. Jang, J. Paek, M. Vieira, R. Govindan, B. Greenstein, A. Joki, D. Estrin, and E. Kohler, "The tenet architecture for tiered sensor networks," in *SenSys*, 2006, pp. 153–166.
- [18] K. Whitehouse, C. Sharp, E. Brewer, and D. Culler, "Hood: a neighborhood abstraction for sensor networks," in *MobiSys*, 2004, pp. 99–110.
- [19] M. Welsh and G. Mainland, "Programming sensor networks using abstract regions," in *NSDI*, 2004, pp. 29–42.
- [20] P. Costa, L. Mottola, A. Murphy, and G. Picco, "Teenytime: transiently shared tuple space middleware for wireless sensor networks," in *Mid-Sens*, 2006, pp. 43–48.
- [21] M. Demirbas and M. Hussain, "A mac layer protocol for priority-based reliable broadcast in wireless ad hoc networks," *BroadNets*, 2006.
- [22] M. Demirbas and S. Balachandran, "Robcast: A singlehop reliable broadcast protocol for wsn," *ADSN*, pp. 1–8, 2007.
- [23] M.-T. Sun, L. Huang, A. Arora, and T.-H. Lai, "Reliable MAC layer multicast in IEEE 802.11 wireless networks," *Proc. ICPP*, pp. 527–537, 2002.
- [24] A. Arora and et. al., "Exscal: Elements of an extreme scale wireless sensor network," *Int. Conf. on Embedded and Real-Time Computing Systems and Applications*, 2005.
- [25] M. Demirbas, O. Soysal, and M. Hussain, "Coordinated locomotion of mobile sensor networks," University at Buffalo, SUNY, Tech. Rep. 2, 2008, under submission.