

H-MAS: A Heterogeneous, Mobile, Ad-hoc Sensor-Network Simulation Environment

Bren C Mochocki Gregory R Madey

University of Notre Dame

Notre Dame, IN 46703

{bmochock, gmadey}@cse.nd.edu

Keywords – mobile, ad-hoc, sensor-network, multi-hop, swarm

ABSTRACT

Ad-hoc Sensor Networks are a feasible and rapidly deployable solution to many common environment monitoring applications. In traditional sensor network design, each unit is capable of sensing, processing and storing its own local data, as well as using a multi-hop forwarding scheme to send this data to off-network storage. Thus, each network element must fill at least four different rolls: sensing, processing, communication, and storage, each of which may be further subdivided. As nano-scale miniaturization becomes feasible, including complete homogeneous functionality at each node will become an increasingly difficult problem. Pico-radio devices provide a platform by which these responsibilities can be offloaded to different nodes, but still remain accessible through wireless communication.

We have developed a conceptual pico-radio sensor network system, and a corresponding design simulation and evaluation environment, H-MAS, that separates the tasks of processing and storage from the sensor nodes. Through agent-based computer simulation using the *Swarm* toolkit, it is possible to compare the effectiveness of current medium access, routing, organization, and energy conservation techniques on heterogeneous, mobile ad hoc sensor networks. Here we provide a proof of concept by deriving several heuristics and design rules for a basic network configuration. The visualization side of H-MAS also provides a convenient way to present the design of MAS systems to non-technical personnel.

1. INTRODUCTION

A **Mobile Ad-hoc Network (MANET)** is a communication network whose topology can change over time, due both to node mobility and limited node energy. Though MANET related research dates back to the early 70's, it has once again come to the forefront of engineering and computer science research. The recent rejuvenation is due to newly developed wireless communication standards (e.g. Bluetooth), a gradual but noticeable paradigm shift from centralized to distributed design (in both control and functionality), and the market's ever increasing demand for mobile computing [5].

The fact that MANETs have no set infrastructure gives them several advantages over conventional networks; namely that the setup time is small, which is ideal for applications such as war zone surveillance/communication, or disaster search-and-rescue missions. MANETs also suffer from several disadvantages, including limited communication range, battery life, limited knowledge of their environment. This lack of infrastructure makes self organization a necessity, and the lack of global knowledge necessitates distributed control policies. Of course, these necessities become marked advantages when the MANET operates in a harsh environment, where node loss is inevitable. Another advantage is that as more elements join the MANET, assuming well designed communication and medium access protocols, the more robust the system becomes. MANETs also lend themselves to less dramatic consumer applications such as inter-vehicle communication for collision avoidance and wireless PDA/Laptop network communication without the benefit of centralized hubs.

A sensor network is in many ways conceptually similar to a MANET. Sensor networks are simply varied arrays of sensors, embedded in or distributed throughout their environment, to collect data over time or detect some event. When compared to MANETs, they themselves are typically more passive in the areas of motion, inter-sensor communication, and ability to influence the environment (though this is not always the case). Proposed applications for sensor networks include equipment supervision, intruder detection, and wildlife observation. Figure 1 illustrates the relationship between sensor networks and MANETs. We will refer to the intersection of sensor nets and MANETs as **Mobile Ad-hoc Sensor nets (MAS)**.




Figure 1. Relation between MANETs and Sensor Nets. Mobile Ad-hoc Sensor Nets fall in both categories.

A number of research groups are working on complete sensor network systems, from the physical to the application layer, including UC-Berkeley's Smart Dust [7], MIT's micro-Adaptive Multi-domain Power aware Sensors (uAmps) [18], and UCLA's Wireless Integrated Sensor Networks (WINS) [19]. These efforts, particularly smart dust, focus on a homogenous distribution of network tasks among nodes.

Though this design methodology may be fine on a larger scale, as devices shrink, it will become more difficult to include all network functionality evenly on every device in the network. For example, the smallest available homogenous MAS design is smart dust from UC Berkeley [7]. The target size for this device is 1-2 mm³, and the smallest actual working implementation is about 7 mm³. Nano-electronic devices, on the other hand, must be on the 1 – 100 nm scale, which is 4 to 6 orders of magnitude smaller. At this level, additional applications will become feasible. One example is to inject a collection of sensors into a human or animal's blood stream to continuously monitor the organism's vital signs, such as immune system activity, blood-sugar level, hormone levels, etc. If these sensors are to reach this small of a scale, all of their parts will inevitably need to be manufactured on the same medium (e.g. silicon, or alternate emerging technology). This leads to realm of System on a Chip (SoC) design.

As described by Bergamaschi [1], SoC design has its own set of problems. For example, process optimizations and voltage supply levels for high performance logic and precision analog designs are often at odds, and the analog devices may be sensitive to nearby switching of digital devices, both of which can increase the design cost and time to market. When coupled with difficulties related to applying differing fabrication techniques to a single device, these costs could reach prohibitive levels. It is unlikely that fabrication techniques for SoC designs in emerging nano-electronic technologies will completely avoid the problems faced by their silicon based ancestors. This is not to say that MAS systems using a homogenous design methodology are impractical, just that employing a heterogeneous methodology offers several clear benefits in both design ease and fabrication cost. Thus, it is likely that future MAS systems, especially on the nano-scale, will be heterogeneous in nature.

Silva et al describe a conceptual MAS system called a PicoRadio network, which consists of a set of nodes that communicate via low power, low range RF transmitters [16]. Though the design targets devices larger than the nano-scale, the system design can inherently handle heterogeneous elements. The design methodology also aggressively works to both minimize system energy consumption and promote energy scavenging; using vibrations and solar power to recharge node batteries. Our simulator uses the PicoNodes

for the basic network element, though we divide the node responsibilities differently (see section 2).

The rest of this abstract is organized as follows: Section 2 reviews the design of the current implementation of H-MAS. Section 3 presents some preliminary experimental results, and section 4 concludes with a look to the future.

2. SIMULATION DESIGN

The purpose of H-MAS is to provide a convenient platform on which to evaluate a variety of MAS configurations at the physical, medium access, network, and application layers, and to extract meaningful design rules from the experimental data. The pros/cons of a heterogeneous vs. homogeneous design can also be explored, as the latter is simply a subset of the former. A secondary design goal is to provide an intuitive visualization that can give insight to the design engineer and casual observer alike.

2.1 Basic Description

The motivational target system is a set of PicoRadio nodes inserted into a stream or lake to continuously measure the state of the environment, such as pH, salinity, or organic material concentration. A parallel concept is a set of sub-PicoRadio nodes (perhaps "FentoRadio" nodes) injected into a human bloodstream to continually measure vital life signs. As such, the mobility of nodes has been defined so they "flow" across the screen on the Y-axis, with some random motion on the X-axis. A future design goal will be to allow the user to create an arbitrary motion description for the nodes. Figure 2 shows several time snapshots of the simulation in progress to illustrate this point. Visually, packet transmission is represented lines. A red line means the packet was received, but rejected because it already exists in the nodes data buffer, and is thus redundant, while white (grey) means it was new, and is inserted into the buffer for the first time.

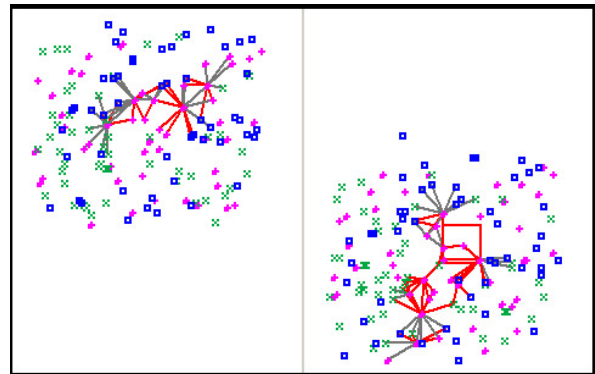


Figure 2. Snapshots of H-MAS in progress. The purple "+" symbols, with various corners filled in, represent different sensors, the green "x" symbols are processors, and the blue boxes are

sinks. Lines are successfully transmitted packets, of which grey ones are accepted, and red ones are rejected. The red box on the right represents a sensor that tried to transmit, but there were no other nodes within its range.

The nodes in the network fall into four basic categories: (1) sensor nodes, (2) processing nodes, (3) sink nodes, and (4) communication nodes. Nodes of type 1 through 3 also fall into category 4. Sensor nodes can be further subdivided by the type of data they measure. Processing nodes fill the role of signal processing, error correction, and data compression. Sinks have the task of storing data locally, and when a large off network data collector becomes available, they forward their stored data to this collector to empty their local storage. This feature has not yet been implemented, but the intended method will be similar to the one presented in [2], modified for a heterogeneous system.

This division of roles differs from the original PicoRadio division in [7]. Kahn et. al divide the network into sensors, actuators, and monitors. Monitors are access points for a user (human, off-network system) to request specific data from the sensors. Actuators are “super nodes” that have a large or unlimited power source, and can somehow influence the environment (e.g. change the air temperature in an office building). Our target application is different from the PicoRadio application in that we are mainly interested in systems that are completely separated from any kind of infrastructure for long periods of time (e.g. when the nodes are flowing down a river in the wilderness) as apposed to consumer applications where some amount of infrastructure is available. Of course, larger actuator nodes could be used in our application as well; to release buffers in the stream to control pH for example. Having super nodes removes the need for individual processors and sinks, because the nodes with an unlimited power source can always take the complete brunt of data processing and storage, unless the super nodes are sparse, and a large amount of processing is required close to the sensing point.

2.2 Design and Implementation

The current implementation was developed in Java using JBuilder 8, on both Windows XP professional and Red-Hat Linux 8.0 with the Swarm toolkit version 2.2 [17]. Experimental data is collected by an instance of a Data collection class that formats the data into an output file and also sends the data to a remote MySQL database.

The class hierarchy, illustrated in Figure 3, cleanly separates the functionality of the physical, medium access, and network layers. Of the listed classes, the current version of H-MAS implements all network node classes, the ObserverSwarm, ModelSwarm, Flooder and GeneralMAC, as well as all associated super classes. The other classes in Figure 3 represent class stubs that will be implemented in

the near future. They include various routing and medium access protocols that will be compared and validated in the H-MAS environment. For a more thorough description and list of these protocols, see section 4. Figure 4 gives more detail about the cardinality, relationships, and navigability between classes.

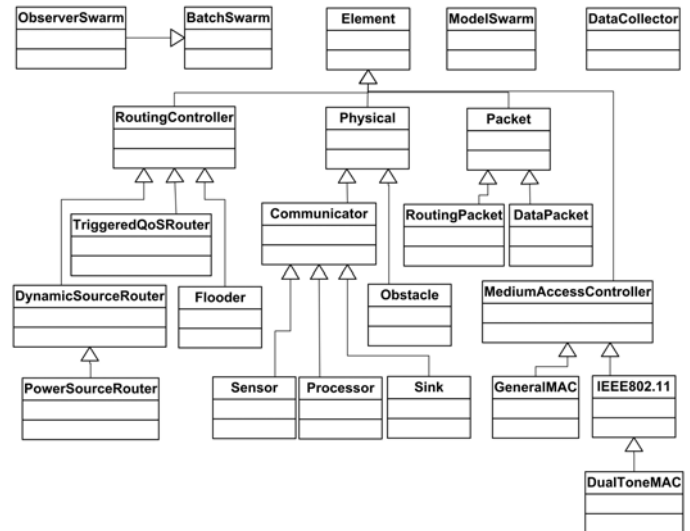


Figure 3. H-MAS class hierarchy.

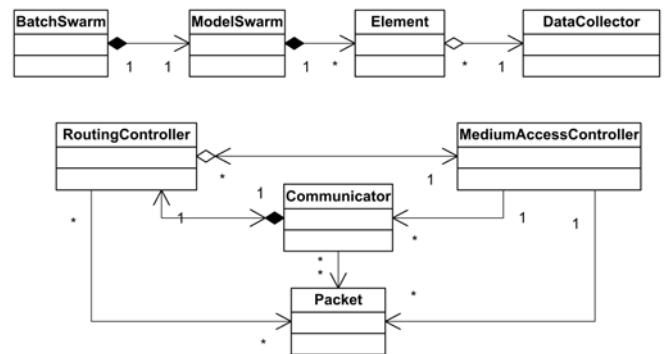


Figure 4. H-MAS class relationships, cardinality and navigability.

Under general medium access, contention for the medium is assumed to never occur, and the only limitation on which nodes can receive a packet is the communicator’s transmission range. One could say that this medium access model is “Ideal” but not realistic. The current protocol used to rout packets is the flooding scheme, similar to the one presented in [8], which is to say that no packet routing is implemented. Addressing is done on the class level. This means that the node which accepts the packet looks at the

packet header, and determines which class the packet is intended for. If the class matches the intended target class, the packet is kept and processed as necessary. If the current node is not of the targeted class, then if its time to live has not expired, it is retransmitted. If the packet is already in the node's received cache, then it is rejected. Packets flow in this way from sensor to sink. Figure 5 illustrates the packet dataflow using this flooding protocol. This simple MAC and routing protocol is used as an initial proof of concept, to make sure that the current design is capable of collect meaning full data, so we can confidently move on to the implementation of more sophisticated protocols.

3. EXPERIMENTAL RESULTS

The initial proof of concept experiment uses an even mix of sensors, processors and sinks, with the flood-based packet communication and ideal medium as described in section 2.

The nodes are randomly distributed to fill 2% of the available spaces in an area that includes 25% of the total (visible) world Y axis and the entire world X axis. This results in exactly 50 of each type of node. Sensor nodes are set to produce a new sensor reading every 1000 time steps, with a random starting time in the range [1, 1000]. The simulation runs for 2000 time steps, which gives every sensor a chance to produce exactly two measurements, for a total of 100 unique sensor packets. Processors spend 10 cycles processing sensor data packets, which is much smaller than the sensor sample frequency. This ensures that processors are almost always free when they receive new data packets. A processor that is currently busy simply forwards the data packet as if it were a simple communicator.

All sensors measured the same type of data, and all sinks were assumed to have infinite local storage. Next, the unique packet miss rate and the number of packets received were measured over the packet time to live (TTL) and node communication range. Figure 6 gives a detailed list of the simulation parameters, Figure 7 plots the miss rate, and Figure 8 plots the number of redundant packets received. Each data point is the average of 10 unique simulation runs, for a total of 1000 simulation runs for the whole experiment.

Given the caveat that this experiment was based on ideal medium conditions with a rather simplistic dataflow, the results illustrated in Figure 7 and Figure 8 are quite interesting. First off, it is clear that the miss rate is very sensitive to the node range, but not the packet TTL, while the number of redundant packets is more sensitive to TTL than node range. From a design perspective, this is very beneficial. We now have two nearly orthogonal parameters with which we can tune system performance to meet whatever design requirements/constraints we may face. For example, if our requirements stated that 90% of unique packets should complete the entire dataflow, and to ensure the robustness of the system to transmission faults we

wanted near 100% redundancy, we could set the range to 8, and the TTL to 5 or 6. Of course these curves will change with more realistic network settings, but the results clearly show the experiment's success as a proof of concept.

4. CONCLUSIONS AND FUTURE WORK

MAS, the intersection of sensor networks and mobile ad-hoc networks, are becoming both an increasingly active research area and more pervasive in the sense of conceptual applications. As more and more real systems are designed on MAS frameworks and nano-electronics make further pervasive applications feasible, it is highly likely that MAS design will move from a homogeneous to heterogeneous design methodology. H-MAS attempts to capture the details of different routing, energy conservation, medium access, and application specific protocols into one complete

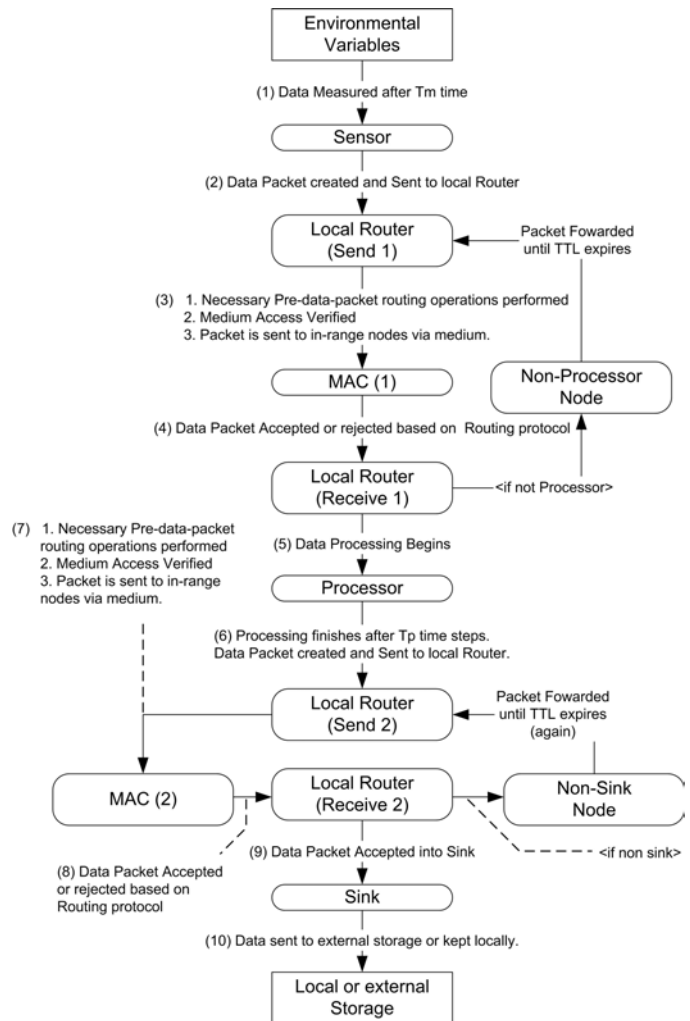


Figure 5. Dataflow of packets through the network with flooding as the routing protocol (i.e. no routing protocol).

simulation environment that is useful to both the system designer when creating the actual product and to the salesman trying to sell the end product to potential customers who lack the technical knowledge to fully grasp detailed design specifications. The experiments on a simplified network in an ideal medium show the usefulness of H-MAS in the former goal, and the visualizations clearly show that latter goal may also be achieved.

```

-----
Simulation Parameters:
-----
NodeStartDensity 0.02
FractionOfWorld 0.25
RouterCount 0
SensorCount 50
ProcessorCount 50
SinkCount 50
MAC 0 (range limited only)
RoutingMethod 0 (flooding)
NodeRange 1 to 10
SensorNodeFraction 0.3333333333333331
ProcessorNodeFraction 0.3333333333333331
SinkNodeFraction 0.3333333333333331
RouterNodeFraction 0.0
TotalFraction 1.0
RunCount 2000
worldXSize 100
worldYSize 300
randomizeUpdateOrder true
RandomSeed -1739848092
SensorMeasurePeriod 1000
DefaultProcessingTime 10
DefaultOutCacheSize 20
DefaultReceiveCacheSize 20
DefaultPacketTimeToLive 1 to 10

```

Figure 6. Simulation parameters.

The proof of concept is complete, but the real work is far from over. First off, we intend to include a detailed energy model into the H-MAS framework that allows for the flexible measurement of energy dissipation at each node, accounts for energy scavenging [7], and implements several power saving strategies, such as the sleep and wakeup methods presented in [7], which uses an ultra low-power wakeup radio, the methods in [14], which requires no special hardware to implement, or those in [9], which dynamically adjust the node transmission range. Once a flexible energy model is implemented that can be easily adjusted to account for various hardware configurations, we can implement more sophisticated routing strategies that are tuned for both a high throughput and low energy consumption, such as Power Source Routing [11], Energy Aware Routing [15], online power aware routing [10], as

well as others [13]. An interesting experiment would be to see how power aware algorithms compare with throughput optimized algorithms, such as dynamic source routing [6], or algorithms optimized for real time applications, such as the Triggered Quality of Service algorithm [3]. Both of these non-power optimized routing algorithms are conceivably useful for systems having a short operation time but require a high level of performance, such as monitoring a patient’s vital signs during an operation, or in networks that are only activated in short lived emergency situations.

Realistic medium access protocols must also be implemented to account for packet collisions and other problems encountered within a non-ideal medium. MAC protocols can use either time-based collision avoidance, such as the IEEE802.11b standard; mixed frequency and time avoidance, such as Dual Busy Tone Multiple Access (DBTMA) [4] or Busy Tone Priority Scheduling (DTPS) [20], or they can be based completely in the frequency domain, like the dynamic channel assignment method presented in [7].

Another important behavior that needs to be included in H-MAS is the ability of sink nodes to efficiently offload stored data to permanent storage, whenever it becomes available. An efficient method designed for the smart dust network is presented in [2], which should be adaptable to the Heterogeneous system in H-MAS.

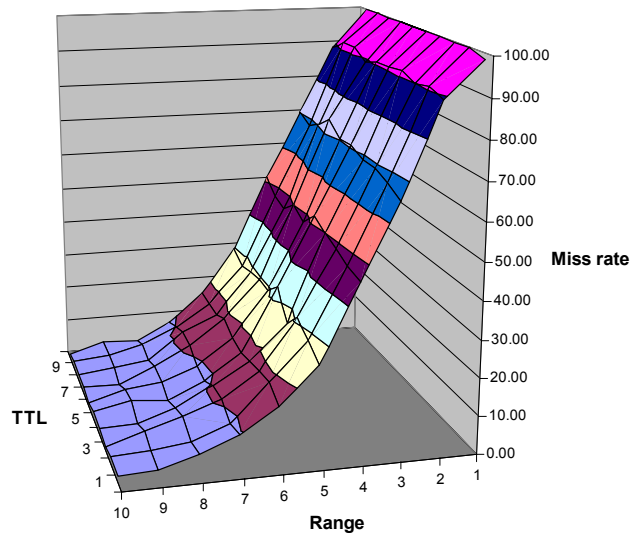


Figure 7. Miss rate of unique data packets vs. packet time to live, and node communication range. The unique packet miss rate is the number of unique sensor packets that did not complete the entire sensor to sink dataflow in Figure 5.




Figure 8. Percent of redundant packets that completed the source to sink dataflow in Figure 5 vs. the packet TTL and node communication range.

The next step to increase the utility of the simulator would be to package it in a convenient, easy to install binary distribution with a clean usable GUI, and tools that make incorporating new protocols and designing new experiments a simple matter. A web based user interface would also allow many researchers and system designers to conveniently run their experiments and share results. A final thought is to extend the simulation to the 3rd dimension, to more accurately model reality.

REFERENCES

- Bergamaschi R. "The A to Z of SoCs" ICCAD, November 10-14, 2002.
- Chatzigiannakis I; Nikolettseas S; Spirakis P. "Smart Dust Protocols for Local Detection and Propagation [Extended Abstract]" POMC'02, October 30-31, 2002.
- De S; Das S; Wu H; Qiao C. "Trigger-Based Distributed QoS Routing in Mobile Ad Hoc Networks" Mobile Computing and Communications Review, Volume 6, No 3.
- Haas Z; Deng J. "Dual Busy Tone Multiple Access (DBTMA): A Multiple Access Control Scheme for Ad Hoc Networks," IEEE Transactions of Communications, 2002
- Haas Z; Deng J; Liang, B; Papadimitratos, P; Sajama, S. "Wireless Ad Hoc Networks." Draft technical report. School of ECE, Cornell University. 2002.
- Johnson D; Maltz D; Hu Y, Jetcheva, J. "The Dynamic Source Routing Protocol for Mobile Ad Hoc Networks (DSR)." IETF MANET Working Group, Internet-Draft. <http://www.ietf.org/ietf/1id-abstracts.txt>, February 21 2002.
- Kahn J; Katz R; Pister K. "Next Century Challenges: Mobile Networking for *Smart Dust*." Mobicom'99 Seattle Washington, 1999.
- Kirkpatrick, M; Madey, G. "Using Swarm intelligence to Rout Messages in Highly Mobile Ad Hoc Networks" Swarm 2003, Seventh annual Swarm meeting, Notre Dame, IN 2003.
- Kubisch M; Karl H; Wolisz A; Zhong L; Rabaey J."Distributed algorithms for transmission power control in wireless sensor networks", IEEE WCNC 2003, New Orleans, LA, March 16-20, 2003.
- Li Q; Aslam J; Rus D. "Online Power-aware Routing in Wireless Ad-hoc Networks." SIGMOBILE'01, July, 2001.
- Maleki M; Dantu K.; Pedram M. "Power-aware Source Routing Protocol for Mobile Ad Hoc Networks." ISLPED'02 August 12-14, 2002.
- Pico-Radio Research group http://bwrc.eecs.berkeley.edu/Research/Pico_Radio/Default.htm, February, 2002
- Safwat A; Hassanein H; Mouftah H. "Energy-Aware Routing in MANETs: Analysis and Enhancements." MSWIM'02, September 28, 2002.
- Schurgers C; Tsiatsis V; Ganeriwal S; Srivastava M. "Optimizing Sensor Networks in the Energy-Latency-Density Design Space." IEEE Transactions on Mobile Computing, Vol. 1, No. 1, January-March 2002.
- Shah R; Rabaey J. "Energy Aware Routing for Low Energy Ad Hoc Sensor Networks", IEEE Wireless Communications and Networking Conference (WCNC), March 17-21, 2002, Orlando, FL.
- Silva J; Shamberger J; Ammer M; Guo C; Li S; Shah R; Tuan T; Sheets M; Rabaey J; Nikolic B; Sangiovanni-Vincentelli A; Wright P. "Design Methodology for PicoRadio Networks" DATE 2001.
- The SWARM development group. <http://www.swarm.org>
- μ -Adaptive Multi-domain Power aware Sensors: <http://www-mtl.mit.edu/research/icsystems/>, February, 2002.
- Wireless Integrated Sensor Networks: <http://www.janet.ucla.edu/WINS/>, February, 2002.
- Yang X; Vaidya N. "Priority Scheduling in Wireless Ad Hoc Networks." MOBIHOC'02, June 9-11, 2002