# ITR/SY: A Distributed Programming Infrastructure for Integrating Smart Sensors

## NSF Program CCR-0121638

**PIs: Umakishore Ramachandran**
Kenneth Mackenzie
Steve DeWeerth
Irfan Essa
Thad Starner

College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280
Phone: (404) 894-5136
FAX: (404) 385-6506
e-mail: rama@cc.gatech.edu
WWW URL: http://www.cc.gatech.edu/~rama

**Annual Report via Fastlane June 30, 2006**

# 1 Activities and Findings

## 1.1 Research and Education

The proposed research is integrating sensing hardware, embedded processing and distributed system support to build a seamless programming infrastructure for ubiquitous presence applications. Fundamental invention and integration of techniques spanning programming idioms and runtime systems for distributed sensors, and building blocks for embedded processing are expected as the primary intellectual contributions of the proposed research. Interfacing these technologies to emerging applications on the one end and novel off-the-shelf sensors at the other end are secondary goals of the proposed research.

In our exploration of novel techniques to support integration of smart sensors into a distributed programming infrastructure, we are conducting focused studies into 1) efficient resource management for streaming applications in wireless sensor networks, cluster computing, and grid computing environments; 2) universal middleware for integrating multiple protocol families; 3) efficient biometric authentication over wireless networks; 4) efficient and robust code dissemination in multihop wireless networks; 5) leveraging web service technologies to support stream based distributed computing applications; and 6) using software (network) caching rather than hardware caching for energy efficient embedded sensor processing. We are also developing specific driving pervasive computing applications to assist in design and evaluation of our distributed infrastructure components. Further, we are developing a sensor lab as a testbed for integrating sensor technologies into pervasive computing applications. This subsection details the research accomplishments this past year (June 2005-May 2006).

### 1.1.1 Distributed Systems Technologies

**Dynamic Data Fusion for Future Sensor Networks.** DFuse [16] is an architectural framework for dynamic application-specified data fusion in sensor networks. It bridges an important abstraction gap for developing advanced fusion applications that takes into account the dynamic nature of applications and sensor networks. Elements of the DFuse architecture include a fusion API, a distributed role assignment algorithm that dynamically adapts the placement of the application task graph on the network, and an abstraction migration facility that aids such dynamic role assignment. Experimental evaluations show that the API has low overhead, and simulation results show that the role assignment algorithm significantly increases the network lifetime over static placement.

**Distributed Garbage Collection Algorithms for Timestamped data.** A journal version of our work with garbage collection issues in Stampede is to appear in IEEE Transactions on Parallel and Distributed Systems [15]. There is an emerging class of interactive multimedia applications that deal with stream data from distributed sources. Indexing the data temporally facilitates ordering individual streams as well as correlating items from different streams. *Stampede* programming system organizes stream data into *channels* that are distributed, synchronized data structures containing timestamped items. A Stampede program is a data flow graph of threads and channels. Stampede semantics for the channels allow concurrent access from multiple threads for input and output. While a channel holds timestamped items, the semantics do not place any restriction on either the order of production or consumption of these items. Further, there is no necessity that the timestamps of the items in a channel be contiguous. These flexibilities are required due to the dynamic and parallel structure of stream-oriented applications targeted by the Stampede system. Under these circumstances, a key issue is the “garbage collection” (GC) of channel items. In this paper we present and compare three different GC algorithms: (1) REF is a simple algorithm that keeps a reference count on individual items; (2) TGC is a distributed algorithm for computing a *global* low water-mark for timestamp values of interest in the entire application; and (3) DGC is another distributed algorithm that uses information about the dependencies between the producers and consumers of data streams to compute a low water mark *local* to each node of the data flow graph. DGC can simultaneously eliminate garbage from channels and unneeded computations from threads. In tests performed using an interactive application, DGC enjoys nearly 30% reduction in the application memory footprint compared to TGC and REF. DGC and REF are also shown to be more scalable compared to TGC.

**MediaBroker++.** MediaBroker is a distributed framework designed to support pervasive computing applications. Key contributions of MediaBroker are efficient and scalable data transport, data stream registration and discovery, an extensible system for data type description, and type-aware data transport that is capable of dynamically transforming data en route from source to sinks. Specifically, the architecture consists of a transport engine and peripheral clients and addresses issues in scalability, data sharing, data transformation and platform heterogeneity.

Continuing from our work reported in the 2003-2004 and 2004-05 annual reports, we have developed *MediaBroker++: Pervasive Computing Middleware for Dynamic Injection, Composition, and Execution of Stream Transformations*. MediaBroker++ is a middleware system that caters to the dynamic needs of applications in a distributed pervasive computing environment. The target environment consists of a wide variety of devices (such as cameras, microphones, and display devices) that act as producers and consumers of stream data, interacting with one another, with the environment, and distributed applications that use them. The MediaBroker++ architecture encompasses several elements. A type server allows clients to dynamically add and remove various data types as well as inject transformation code that operates on the streams 0owing through the system. Upon request from a client, the system dynamically composes the set of transforms (orchestrating the plumbing to get the streams to the transforms) to satisfy the request. The system can handle multiple requests simultaneously. A transformation engine serves as a logical entity that is responsible for the creation and management of the execution environments for the transforms (called transformation environments) on the available computing resources. A scheduler, in concert with the transformation engine, executes the transformations as they become ready to run on the available resources. Upon dispatch by the scheduler, each transformation environment, working independently, /rst retrieves the corresponding transformation code from the type server and then computes and returns the result, possibly enabling other transforms to become ready for execution. The scheduler generates tree-structured query plans, utilizes a post-order tree serialization algorithm, and monitors and assigns the ready transforms to the available transformation environments for execution.

MediaBroker++ [18] has garnered a significant amount of attention from the research community. In particular, there is a strong interest in collaborating with UCI's RESCUE project to use MediaBroker++ as the underlying system infrastructure for the application oriented research being pursued in that project.

**SensorStack.** Wireless Sensor Networks are deployed in demanding environments, where application requirements as well as network conditions may change dynamically. Thus the protocol stack in each node of the sensor network has to be able to adapt to these changing conditions. Historically, protocol stacks have been designed with strict layering and strong interface between the layers leading to a robust design. However, crosslayer information sharing could help the protocol modules to make informed decisions and adapt to changing environmental conditions. There have been ad hoc approaches to facilitating cross-layer cooperation for adaptability. However, there has been no concerted effort at providing a uniform framework for crosslayer adaptability that preserves the modularity of a conventional protocol stack. We present a novel service, information exchange service (IES), as a framework for cross-module information exchange. IES is a centrally controlled bulletin-board where different modules can post available data, or request for useful information, and get notified when the information becomes available. IES is integrated into a proposed SensorStack architecture that preserves the stackability of the protocol layers while facilitating adaptability. IES has been implemented in TinyOS and Linux, to show both the feasibility of the design as well as demonstrate the utility of cross-layering to increase application longevity. *SensorStack* [7] is a research we have undertaken to provide system support for cross-layering in sensor network stack.

A provisional patent has been filed for this research as well [6].

**Packet Delivery in Sensor Networks.** The PI, working with Professor Mostafa Ammar and a student has been exploring issues with respect to the reliability of packet delivery in dense wireless sensor networks [20].

Wireless sensor networks (WSN) built using current Berkeley Mica motes exhibit low reliability for packet delivery. There is anecdotal evidence of poor packet delivery rates from several field trials of WSN deployment. All-to-one communication pattern is a dominant one in many such deployments. As we scale up the size of the network and the traffic density in this communication pattern, improving the reliability of packet delivery performance becomes very important. This study is aimed at two things. Firstly, it aims to understand the factors limiting reliable packet delivery for all-to-one communication pattern in dense wireless sensor networks. Secondly, it aims to suggest enhancements to well-known protocols that may help boost the performance to acceptable levels. We pick three protocols, namely, Flooding, AODV, and Geographic routing as candidates for this study. We first postulate the potential reasons hampering packet delivery rates with current CSMA-basedMAC layer used by the radios deployed in WSN. We then propose a set of enhancements that are aimed to mitigate the ill-effects of these factors. Using TOSSIM, we perform a detailed study of these protocols and the proposed enhancements. This study serves several purposes. First, it helps us to quantify the detrimental effects of these factors. Second, it helps us to quantify the extent to which our proposed enhancements improves packet delivery performance. Concretely, we show that using Geographic routing in a WSN with 225 nodes spread over 150 feet x 150 feet, the proposed enhancements yield a 23-fold improvement in packet delivery performance over the baseline. Further, the enhancements result in fairness (measured by the number of messages received from each node at the destination). Lastly, we show that the overhead (in terms of retransmissions, acknowledgement messages, and control messages) is reasonable.

### 1.1.2 Sensor Technologies: Heterogeneous Wireless Sensor Network

We have created a heterogeneous wireless sensor network (HWSN) testbed consisting of Berkeley Mica-2 motes, skiffs, iPAQs, and Stargates. The Stargate nodes serve as a gateway between the low bandwidth network (Chipcon radio) consisting of the Mica-2 motes, and the high bandwidth network (802.11) consisting of the skiff nodes and the iPAQs.

This testbed serves as a platform for experimenting with applications such as video-based surveillance using the distributed systems technologies we have developed in this project. More importantly, we wish to solve system issues and network protocol issues for HWSN.

### 1.1.3 Sensor Technologies: RFID Testbed

We are also creating an RFID testbed. We will use this testbed to explore system support for reliable RFID deployment. The reliability of RFID systems depends on a number of factors including RF interference, deployment environment, configuration of the readers, and placement of the readers and the tags. While the RFID technology is improving rapidly, a reliable deployment of this technology is still a significant challenge thus impeding a wide-spread adoption.

Our research is focused on evolving system software solutions for achieving highly reliable deployment that takes into account all the sources of inherent unreliability in the underlying RFID technology. The output from this research is both to serve a diagnostic and a prescriptive role. The deliverables from our research to a system administrator are two-fold: (1) diagnosis of the sources of unreliability in an existing deployment; and (2) prescription for a new deployment given the environmental conditions.

### 1.1.4 Textbook Development

In previous reports, we have discussed a course the PI (Ramachandran) has developed and teaches, CS 2200: An introduction to systems and networks. This is a novel approach to teaching a first course in systems in an integrated fashion combining the hardware and software issues in one introductory course. Due to its novelty, there is no textbook that serves the purpose for this course. In Spring 2005, the PI (Ramachandran) wrote a comprehensive set of notes for this course. In Spring 2006, he has signed a contract with Addison-Wesley to have the book published as a textbook [17].

## 1.2 Training and Development

Needless to say, graduate students pursuing their doctorate are the primary focus of training and development. Recent graduates from our group partially funded by this ITR grant include Matt Wolenetz (Microsoft, Redmond), Sameer Adhikari (Intel, Portland), Zack Kurmas (Grand Valley State University, Michigan), Arnab Paul (Intel, Portland), and Josh Fryman (Intel, Portland). Four students (Nissim Harel, Namgeun Jeong, Bikash Agarwalla, and Rajnish Kumar) are in the final stages of their dissertation research and are expected to graduate in August/September timeframe.

The PI (Ramachandran) has been involving his senior graduate students in developing ideas and partially writing proposals for funding agencies such as NSF to prepare them for the road ahead when they graduate. He also tried a new experiment these past two years. For the conferences (such as PerCom 2005 and 2006) on which he served as a program committee member, he created a mock program committee with his graduate students. One of them (usually the senior most graduate student) serves as the PC chair; distributes the papers to the graduate students for review; collects the feedback and conducts a mock PC meeting going over the papers, rank ordering them, and making recommendations of accept/reject. Of course, I supervise the whole process. This has been a great hit with his graduate students since it gives them a first hand experience of how such meetings are run and how decisions are made in conference program committees.

We continue to attract bright and interested graduates and undergraduates to research projects in our group. Undergraduate participation in research within the College is facilitated by the excellent UROC program (`www.cc.gatech.edu/program/uroc`), coordinated by Professor Amy Bruckman. A variety of institute-wide programs are also available (`www.undergraduateresearch.gatech.edu`) including a special fund sponsored by the president of Georgia Tech (PURA) and several NSF-sponsored projects. We were pleased to support several undergraduates on our ITR-related projects during Summer/Fall 2005 and Spring/Summer 2006. They were: Ryan Juang (MediaBroker), Seth Horrigan (MediaBroker), Gabriel Heim (Heterogeneous Wireless Sensor Networks), Charles Reiss (Heterogeneous Wireless Sensor Networks), Puyan Lotfi (Heterogeneous Wireless Sensor Networks), Scott Gilliland (Heterogeneous Wireless Sensor Networks), Paolo Rugero Mentonelli (System support for RFID), Echezona C. Ukah (System support for RFID), and David Worsham (System support for RFID). For details of the PURA program, along with a list of recipients, see the website.

Many of the ongoing ITR-related projects are partially staffed by students working in the context of the Systems Hackfest. This is a group of undergraduates who participate in various research projects for pay, course credit, or just for fun. Hackfest is supervised by Research Scientist Phil Hutto and runs throughout the year. Summer sessions are most productive and have recently involved 6-10 students. Students meet briefly in a weekly session to report progress and plan milestones for the coming week. The group meeting allows cross-fertilization of project ideas and helps to educate the students. In addition, it provides an opportunity for group brainstorming on design and debugging issues. Weekly project meetings are focused on specific research tasks and often involve relevant faculty, grad students and staff.

During the last year undergraduates have participated in the following projects: MediaBroker, Heterogeneous Wireless Sensor Networks, and RFID system support. We are also pleased by the number of undergraduates in our group who continue on to graduate study both here at Georgia Tech and at other top schools.

We believe the Hackfest is an excellent opportunity for initiating undergraduates into the form and substance of academic research. In addition, the size and maturity of the inter-related research efforts provides a fertile matrix for varied interactions and training. Each group – undergraduates, Masters students, PhD students, research scientists and senior faculty – have regular opportunities for cross-group interactions. For example, undergraduates can look to senior faculty for vision and research goals, to research scientists for design advice, to graduate students for technical assistance and literature questions, and to each other for day to day camaraderie.

## 1.3 Outreach

Through the auspices of the Center for Experimental Research in Computer Systems (CERCS) we continue to invite and host key individuals from academia and commercial research labs engaged in complementary research. We have an industrial advisory board comprising of researchers from all the leading industries. We have two board meetings every year (one in the Fall is in conjunction with the NSF/IUCRC annual meeting), and one of the fixed items on the agenda during these meetings is a poster and demo session by the students. This serves as a great opportunity for students (graduate and undergraduate) working on this award to meet and discuss their research with leading researchers from industries.

### 1.3.1 Inter-university Collaboration

The PI (Ramachandran) has been continually exploring opportunity to export the technologies developed in this grant to other institutions to increase the impact of this work. Currently, we are actively exploring the following opportunity: The ITR RESCUE project (`www.itr-rescue.org`) has goals very similar to our ITR project but comes from an application perspective and is predicated on the availability of a distributed programming infrastructure for its success. Since the latter is exactly the focus of our ITR it seems like an opportunity for a fruitful partnership. We have been discussing with Professors Sharad Mehrotra (PI on the RESCUE project), Ramesh Jain, and Nalini Venkatasubramanian of UCI, plans to make this opportunity real by exchanging students and/or research scientists to put their applications on top of our distributed systems infrastructure.

### 1.3.2 Technology Transfer

The Stampede programming systems developed under the auspices of a companion NSF award is being used as the basis for a dynamic cluster scheduling framework for use in technical applications at the Federal Reserve Bank - Atlanta. Two PhD students, Namgeun Jeong and David Hilley, are working at FRB-Atlanta as year-round interns helping the economists at FRB to run their economic forecasting applications on high-performance clusters using this dynamic scheduling framework.

The Stampede programming model also served as an inspiration for the T-Streams programming model that Dr. Kath Knobe developed at HP Cambridge Research Lab. T-Streams is being used for supporting scientific applications on high performance clusters.

The DFuse fusion architecture developed under this award is receiving attention from different sources as a potential target for commercialization (`www.dodtechmatch.com/DOD/Opportunities/SBIRView.-aspx?id=A06-112`). We are partnering with Reservoir Labs (`www.reservoir.com`) to explore such an opportunity.

# 2 Publications and Products

## 2.1 Publications

See the references at the end of this document for publications that appeared in the period that covers this annual report.

## 2.2 Web Site

Please visit the project web site at www.cc.gatech.edu/~rama/ubiq-presence/

# 3 Contributions

The activities we are currently undertaking have resulted in a significant number of publications and software artifacts. These are listed in the references at the end of this report.

## 3.1 Human Resources

Roughly 10 graduate students, 10 undergraduate students, 2 research scientists, and 1 visiting faculty member have been associated with this project during the period covered by this annual report.

## 3.2 Student Placement

We continue to place student members of our research group in interesting project-related internships, graduate programs and industry jobs.

Undergraduate alumni of this project are now grad students at the following places:

- Gabriel Heim is starting his graduate work at Georgia Tech from Fall 2006.
- David Hilley is a PhD student working under Professor Ramachandran's tutelage as of Fall 2005.
- Ilya Bagrak is currently a graduate student at University of California, Berkeley. He received an honorable mention by the CRA for outstanding undergraduate and received an NSF and DOE graduate fellowship in 2004.
- Ansley Post is a graduate student at Rice University.
- Kirill Mechitov is a graduate student at UIUC.

Several students trained by this projects are sought after for pursuing internships in industries:

- PhD student Xiang Song spent the Summer and Fall of 2005 working with Dr. Raj Kumar at HP Labs in Palo Alto on grid infrastructure for "appliance computing." This is a continuation of the work he did for that group in Summer/Fall 2004.
- PhD student Rajnish Kumar interned at NEC labs.
- PhD student Dave Lillethun interned at Motorola labs.
- Professor Ramachandran has an ongoing relationship with the Federal Reserve Bank of Atlanta under which two graduate students are employed as interns year-round.
  PhD student Namgeun Jeong has been interning with them since January 2004.
  PhD student Hasnain Mandviwala worked since January 2004 until December 2005.
  PhD student David Hilley worked from January 2006 to May 2006.
  Some of the technologies developed through partial support from the ITR grant such as the Stampede system is being used as a dynamic cluster scheduling framework for running compute intensive applications developed by economists.
- PhD student Hasnain Mandviwala has been interning with Intel research since January 2006 to carry out collaborative research of the PI (Ramachandran) with Dr. Kath Knobe of Intel.

Recent alumni of this project include:

- Undergraduate Zachary Crowell finished his degree in Summer 2005 and joined Microsoft (this was enabled by an internship he did in the Summer of 2004);
- Yavor Angelov is pursuing his PhD under Professor Ramachandran while employed at Microsoft.
- Dr. Zachary Kurmas [9] graduated in December 2004, and has taken up the position of Assistant

Professor, Department of Computer Science & Information Systems, Grand Valley State University, Michigan.

- Dr. Sameer Adhikari [1] graduated in December 2004 and joined Intel, Portland as a system engineer.
- Dr. Arnab Paul [13] graduated in May 2005 and has joined Intel, Portland.
- Dr. Josh Fryman [4] graduated in August 2005 and has joined Intel, Santa Clara.
- Dr. Matt Wolenetz [22] graduated in August 2005 and has joined Microsoft, Redmond.

### 3.3 Research and Education

The research artifacts from the project are finding their way into graduate courses and we have significant undergraduate participation in project-related research. We have funded a number of undergraduates through the REU supplement attached to this ITR grant, sponsored a number of independent undergraduate research projects for course credit (CS 4903), and have sponsored capstone senior design projects (CS 3901) that each result in a poster presentation at the annual Undergraduate Research Symposium.

The ITR project has reinforced the connectedness of hardware and software and the need to train students in system architecture quite early in their undergraduate preparation. With this in mind, Professor Ramachandran embarked on writing a set of course notes [17] for use in the sophomore level course on systems and networks. The course notes were well received by the students. notes for this course. In Spring 2006, he has signed a contract with Addison-Wesley to have the book published as a textbook.

## 4 Special Requirements

The total budget for this 5-year proposal is $1.35M. However, due to fiscal constraints NSF chose to front-load the total budget by awarding $750,000 in the first year. The second, third, and fourth year increments were $200,000 each. The understanding with the program manager (Dr. Helen Gill) is that our spending plan for the award will be more balanced over the 5-year period despite the front-loaded nature of the allocation.

Further, due to the recent collaboration plans with UCI, we have requested and received an informal approval for expending the remaining funds in the project by August 2008.

## References


[1] Sameer Adhikari. *Programming Idioms for Pervasive Computing*. PhD thesis, College of Computing, Georgia Institute of Technology, December 2004.

[2] Bikash Agarwalla, Nova Ahmed, David Hilley, and Umakishore Ramachandran. Streamline: A scheduling heuristic for streaming applications on the grid. In *Thirteenth Annual MultiMedia Computing and Networking Conference (MMCN 2006)*, January 2006.

[3] Yavor Angelov, Umakishore Ramachandran, Kenneth Mackenzie, James Matthew Rehg, and Irfan Essa. Experiences with optimizing two stream-based applications for cluster execution. *Journal of Parallel and Distributed Systems*, June 2005.

[4] Josh Fryman. *SoftCache Architecture*. PhD thesis, College of Computing, Georgia Institute of Technology, August 2005.

[5] Rajendra Kumar and Xiang Song. Gridlite: An infrastructure for provisioning and managing grid services in resource constrained devices. Abstract appeared in GlobusWorld 2005. http://www.globusconsortium.org/journal/20050907/gridlite.html; also available as technical report http://www.hpl.hp.com/techreports/2005/HPL-2005-146.html.

[6] Rajnish Kumar, Santashil PalChaudhuri, and Umakishore Ramachandran. An adaptable protocol stack architecture for wireless sensor network, June 2006. Provisional patent application for GTRC ID No. 3710.

[7] Rajnish Kumar, Santashil PalChaudhuri, and Umakishore Ramachandran. System support for cross-layering in sensor network stack, 2006. Unpublished manuscript, submitted for publication to a conference.

[8] Rajnish Kumar, Arnab Paul, and Umakishore Ramachandran. Fountain broadcast for wireless sensor networks. In *Proceedings of the Second International Workshop on Networked Sensing Systems (INSS 2005)*, San Diego, CA, June 2005.

[9] Zachary Kurmas. *Generating and Analyzing Synthetic Workloads using Iterative Distillation*. PhD thesis, College of Computing, Georgia Institute of Technology, December 2004.

[10] Martin Modahl, Bikash Agarwalla, T. Scott Saponas, Gregory Abowd, and Umakishore Ramachandran. Ubiqstack: a taxonomy for a ubiquitous computing software stack. *Personal and Ubiquitous Computing*, 10(1):21–27, February 2006.

[11] Jin Nakazawa, W. Keith Edwards, Hideyuki Tokuda, and Umakishore Ramachandran. A bridging framework for universal interoperability in pervasive systems. In *26th International Conference on Distributed Computing Systems (ICDCS 2006)*, 2006. To Appear.

[12] Jin Nakazawa, Hideyuki Tokuda, W. Keith Edwards, and Umakishore Ramachandran. $g^2ui$: A geographical user interface for multimedia exploration using diverse off-the-shelf devices, 2006. Unpublished manuscript, submitted for publication to a conference.

[13] Arnab Paul. *Application of Error Correcting Codes to Distributed Pervasive Computing*. PhD thesis, College of Computing, Georgia Institute of Technology, December 2004.

[14] Arnab Paul, Rajnish Kumar, and Umakishore Ramachandran. LAWN: A lightweight authentication protocol for wireless networks. In *Proceedings of the 4th IEEE International Conference on Network Computing Applications (NCA 05)*, July 2005.

[15] Umakishore Ramachandran, Kathleen Knobe, Hasnain A. Mandviwala, and Nissim Harel. Distributed garbage collection for timestamped data. *IEEE Transactions on Parallel and Distributed Systems*, 2006. To Appear.

[16] Umakishore Ramachandran, Rajnish Kumar, Matthew Wolenetz, Brian Cooper, Bikash Agarwalla, JunSuk Shin, Phillip Hutto, and Arnab Paul. Dynamic data fusion for future sensor networks. *IEEE Transactions on Sensor Networks*, 2006. To Appear.

[17] Umakishore Ramachandran and Bill Leahy. Course notes for CS 2200 - introduction to systems and networks, 2005. Draft of textbook for CS 2200, URL: http://swiki.cc.gatech.edu:8080/cs2200/1.

[18] Umakishore Ramachandran, David Lillethun, Bin Liu, Jin Nakazawa, David Hilley, Seth Horrigan, and Brian Cooper. Mediabroker++: Pervasive computing middleware for dynamic injection, composition, and execution of stream transformations, 2006. Unpublished manuscript, submitted for publication to a conference.

[19] Umakishore Ramachandran, Martin Modahl, Ilya Bagrak, Matthew Wolenetz, David Lillethun, Bin Liu, James Kim, Phillip Hutto, and Ramesh Jain. MediaBroker: a pervasive computing infrastructure for adaptive transformation and sharing of stream data. *Elsevier Journal on Pervasive and Mobile Computing, special issue*, 1(2):257–276, July 2005.

[20] Junsuk Shin, Mostafa Ammar, and Umakishore Ramachandran. On improving the reliability of packet delivery in dense wireless sensor networks, 2006. Unpublished manuscript, submitted for publication to a conference.

[21] Xiang Song and Rajendra Kumar. Gridlite: An infrastructure for provisioning and managing grid services in resource constrained devices. In *Proceedings of The 2006 International Conference on Grid Computing & Applications*, pages 127–133. CSREA Press, 2006.

[22] Matthew Wolenetz. *Characterizing Middleware Mechanisms for Future Sensor Networks*. PhD thesis, College of Computing, Georgia Institute of Technology, August 2005.

[23] Matthew Wolenetz, Rajnish Kumar, Junsuk Shin, and Umakishore Ramachandran. A simulation-based study of wireless sensor network middleware. *International Journal on Network Management, special issue on sensor networks*, 2005.