Final Report for Period: 09/2007 - 08/2008 **Submitted on:** 12/05/2008

Principal Investigator: Ramachandran, Umakishore . **Award ID:** 0330639

Organization: GA Tech Res Corp - GIT

Submitted By:

Ramachandran, Umakishore - Principal Investigator

Title:

NMI: Exploration of Middleware Technologies for Ubiquitous Computing with Applications to Grid Computing

Project Participants

Senior Personnel

Name: Ramachandran, Umakishore

Worked for more than 160 Hours: Yes

Contribution to Project:

Name: Abowd, Gregory

Worked for more than 160 Hours: Yes

Contribution to Project:

Name: Wolenetz, Matt

Worked for more than 160 Hours: Yes

Contribution to Project:

Research scientist, carrying out research on STAGES, DFuse, EventWeb, (Federated) MediaBroker, Family Intercom, and TV Watcher. Matt is actively involved in most aspects of the project and is focusing on DFuse related work for his dissertation.

Name: Hutto, Phil

Worked for more than 160 Hours: No

Contribution to Project:

Research scientists overseeing resource acquisition, undergraduate student guidance.

Name: Kumar, Raj

Worked for more than 160 Hours: No

Contribution to Project:

Industry collaborator hosting interactive grid research internships at HP Labs, Palo Alto, CA.

Name: Basu, Sujoy

Worked for more than 160 Hours: No

Contribution to Project:

Industry collaborator hosting interactive grid research internships at HP Labs, Palo Alto, CA.

Name: Starner, Thad

Worked for more than 160 Hours: No

Contribution to Project:

Prof. Starner is one of the faculty directly involved with the TVWatcher application research, collaborating on correlation techniques and usability studies of the application.

Name: Essa, Irfan

Worked for more than 160 Hours: No

Contribution to Project:

Prof. Essa is one of the faculty directly involved in the TVWatcher application project, and collaborates on vision-based processing research used in the project for feature extraction and correlation.

Name: Yu, HeonChung

Worked for more than 160 Hours: No

Contribution to Project:

Prof. Yu is on the faculty at Korea University in Seoul, S. Korea. He is on sabbatical for a year starting January 2004 and is spending this time in residence at the College of Computing at Georgia Tech. His research area is fault-tolerance and resource scheduling in grid systems. He is working with us on our efforts to dynamically support offloading of computationally intensive sensor processing on a computation grid.

Name: Rehg, James

Worked for more than 160 Hours: No

Contribution to Project:

Prof. Rehg is one of the co-authors of the original Stampede system (while at Compaq CRL) and continues to work with us on semantically rich processing of visual sensor data. He also works on sensor fusion issues.

Name: Jain, Ramesh

Worked for more than 160 Hours: No

Contribution to Project:

Prof. Jain leads the EventWeb project that is one of the driving applications for our MediaBroker work and our time/space/location reasoning system. He has considerable expertise in media processing and automated 'event' detection. Prof. Jain holds joint appointments in CoC and ECE and holds the Rhesa S. Farmer, Jr. Distinguished Chair in Embedded Experiential Systems within the College of Computing.

Name: Edwards, Keith

Worked for more than 160 Hours: Yes

Contribution to Project:

Professor Edwards was a collaborator on the uMiddle project and the G2UI project with the PI Ramachandran and the visitor Dr. Jin Nakazawa. He was not supported by this award.

Post-doc

Name: Singh, Rahul

Worked for more than 160 Hours: No

Contribution to Project:

Dr. Singh is a post-doc working with Ramesh Jain. He has been coordinating the EventWeb project from Prof. Jain's side.

Name: Kumar, Rajnish

Worked for more than 160 Hours: No

Contribution to Project:

Rajnish completed his PhD in Fall 2006. He is continuing with me on another NSF grant as a research scientist.

Graduate Student

Name: Agarwalla, Bikash

Worked for more than 160 Hours: Yes

Contribution to Project:

Bikash is one of the primary developers and maintainers of Stampede and D-Stampede. He is involved in a variety of aspects of the project, especially work on grid technologies for supporting pervasive computing applications. He spent 9 months working at HP Labs in Palo Alto. Recently, he has concentrated on federating the Stampede runtime and on researching requirements for interactive grid schedulers.

Name: Bian, Xuehai

Worked for more than 160 Hours: Yes

Contribution to Project:

Name: Yang, Kevin

Worked for more than 160 Hours: Yes

Contribution to Project:

Kevin worked on a vision infrastructure for locationing in the Aware Home.

Name: Modahl, Martin

Worked for more than 160 Hours: Yes

Contribution to Project:

Martin worked with the MediaBroker project with funding from PURA award from Georgia Tech as an undergraduate. He continued and completed the Masters program, and has been working on EventWeb (with federated MediaBroker as base).

Name: Lillethun, Dave

Worked for more than 160 Hours: No

Contribution to Project:

Dave has recently joined our group as an incoming graduate student. He is working to concretize a Federated MediaBroker implementation that will use Federated Stampede and our time/space/identity reasoning system to provide the underlying runtime for the EventWeb application.

Name: Ahmed, Nova

Worked for more than 160 Hours: No

Contribution to Project:

Nova is a graduate student in Computer Science at Georgia State University. She is working on high-level scheduling of grid-based computational processing of sensor data.

Name: Edwards, Ken

Worked for more than 160 Hours: No

Contribution to Project:

Ken worked with us as an undergraduate on the TVWatcher project and is entering the Master's program in the College of Computing at Georgia Tech this semester. He has been working on resource monitoring and the information subsystem of the grid infrastructure in service of our grid scheduling research.

Name: Hilley, David

Worked for more than 160 Hours: No

Contribution to Project:

David is a primary contributor to the TVWatcher project's current implementation and evaluation. Beginning as an undergraduate researcher, he worked on an 'On-line Television Channel Recommender' project. He also received funding from PURA award from Georgia Tech as an undergraduate. He has continued into the Masters program and continues extensive work on the project, more recently with specialized video feature extraction algorithms under joint supervision of Prof. Essa, Prof. Ramachandran and Prof. Starner.

Name: El-Helw, Ahmed

Worked for more than 160 Hours: No

Contribution to Project:

Ahmed worked as an undergraduate, with funding from PURA award from Georgia Tech, on the 'On-line Television Channel Recommender' project, now named TVWatcher. Ahmed implemented the text correlation engine, web browser pop-ups, and web service interfaces. He has entered the Masters program and continues work on the project.

Name: Song, Xiang

Worked for more than 160 Hours: No

Contribution to Project:

Xiang Song is spending the Summer and Fall of 2004 working with Dr. Raj Kumar at HP Labs in Palo Alto on grid infrastructure for 'appliance computing.'

Name: Pack, Derik

Worked for more than 160 Hours: No

Contribution to Project:

Derick just completed his Master's thesis on an architecture for EventWeb. He has done a large amount of implementation and integration of the current prototype system.

Name: Liu, Bin

Worked for more than 160 Hours: No

Contribution to Project:

Bin is a PhD student working with Prof. Jain's EventWeb group. He is exploring streaming database issues such as a language for continuous queries, and synchronization and query optimization.

Name: Wei, JinPeng

Worked for more than 160 Hours: No

Contribution to Project:

JinPeng completed a 'mini-project' related to the DFuse work and was supervised by PhD students Rajnish Kumar and Matthew Wolenetz. This project was the primary outcome of a special 'research orientation' (CS7001) course required of all incoming PhD students in the College of Computing at Georgia Tech.

Undergraduate Student

Name: Saponas, Scott

Worked for more than 160 Hours: Yes

Contribution to Project:

Scott worked with Prof. Abowd to develop a more flexible Digital Family Portrait implementation based upon a UPnP-mapped locationing subsystem.

Name: Kim, James

Worked for more than 160 Hours: No

Contribution to Project:

James was part of a group of students developing Java-based Sensor Lab infrastructure during the summer of 2004. He developed a Mediator for the ER1 mobile robot, enabling applications to use Sensor Lab services to interact with the robot. In addition, James is continuing the implementation of a Federated MediaBroker (begun by Martin Modahl) for his Senior Design project in Fall 2004.

Name: Fletcher, Jason

Worked for more than 160 Hours: No

Contribution to Project:

Jason is beginning research with our group as an undergraduate in Fall 2004, working with the TVWatcher group to develop audio-based feature extractors to enhance the level of information available to the correlation engine.

Name: Bagrak, Ilya

Worked for more than 160 Hours: Yes

Contribution to Project:

Ilya worked primarily with the MediaBroker project as an undergraduate with funding from PURA award from Georgia Tech, designing the core of the type-map based data transformation engine. Ilya won CoC Outstanding Undergraduate Research Assistant and was awarded Honorable Mention in the CRA nomination. He now attends graduate school at Berkeley.

Name: Crowell, Zachary

Worked for more than 160 Hours: No

Contribution to Project:

Zachary has worked on the EventWeb project for a semester and the TVWatcher project for another semester. He also developed the GUI-based Intercom Client component for the Family Intercom application, integrated with a UPnP-based status and control framework.

Name: Thomas, Robert

Worked for more than 160 Hours: No

Contribution to Project:

Robert was part of a group of students developing Java-based Sensor Lab infrastructure during the summer of 2004 and will continue into the Fall of 2004. Robert developed a Mediator for the Versus Tech RF/IR badge tracking system deployed within the College. In addition, Robert developed a Mediator for scrolling marquee signs and developed the Registry used for device discovery and for download of Mediator proxies (bytecodes). These Mediators enable applications to interact with Sensor Lab devices.

Name: Horrigan, Seth

Worked for more than 160 Hours: No

Contribution to Project:

Designed and implemented the tye server for MB++

Name: Young, Sam

Worked for more than 160 Hours: No

Contribution to Project:

Helped in interfacing MobiGo to uMiddle

Name: Slaughter, Michael

Worked for more than 160 Hours: No

Contribution to Project:

UI issues in the RF2ID project

Name: McCauley, Matthew

Worked for more than 160 Hours: No

Contribution to Project:

Application development for camera sensor network

Name: Woo, Joseph

Worked for more than 160 Hours: No

Contribution to Project:

Application development for camera sensor network

Name: Parker, Robert

Worked for more than 160 Hours: No

Contribution to Project:

Application development for camera sensor network

Technician, Programmer

Name: Brennan, Sean

Worked for more than 160 Hours: Yes

Contribution to Project:

Sean is a staff member of the campus computing group (OIT) at Georgia Tech. He has been assigned part-time to work on a prototype of the EventWeb system featuring campus seminars and other events at Georgia Tech. He has done extensive work provisioning and setting up specialized facilities on campus to accomodate the EventWeb system.

Other Participant

Name: Yip, ManFung

Worked for more than 160 Hours: No

Contribution to Project:

Active RFID and path planning with robots mounted with RFID reader

Research Experience for Undergraduates

Organizational Partners

HEWLETT PACKARD

Dr. Raj Kumar, is a Co-PI on this proposal (with no monetary commitment to NSF).

My students routinely spend extended periods of time working with Dr. Raj Kumar's group. The most recent examples (Bikash Agarwalla and Xiang Song).

Keio University
Jin Nakazawa visited us at Georgia Tech from Keio.

Other Collaborators or Contacts

Activities and Findings

Research and Education Activities:
See attached PDF file (Section Activities and Findings)

Findings:
See attached PDF file (Section Activities and Findings)

Training and Development:
See attached PDF file (Section Training and Development)

Outreach Activities:
See attached PDF file (Section Outreach)

Journal Publications

Books or Other One-time Publications

Web/Internet Site

URL(s):
http://www.cc.gatech.edu/~rama/nsfnmi
Description:

Other Specific Products

Contributions

Contributions within Discipline:
See attached PDF file (Section Contributions)

Contributions to Other Disciplines:

Contributions to Human Resource Development:
See attached PDF file (Section Contributions)
Contributions to Resources for Research and Education:
See attached PDF file (Section Contributions)
Contributions Beyond Science and Engineering:

Conference Proceedings

Categories for which nothing is reported:

Any Journal
Any Book
Any Product
Contributions: To Any Other Disciplines
Contributions: To Any Beyond Science and Engineering
Any Conference

Exploration of Middleware Technologies for Ubiquitous Computing with Applications to Grid Computing

NSF Middleware Initiative 03-513, Award 0330639

PIs:
Umakishore Ramachandran
Gregory Abowd
Senior Personnel:
Dr. Raj Kumar*
Dr. Sujoy Basu*

College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280
Phone: (404) 894-5136
FAX: (404) 385-2295
e-mail: {rama, abowd}@cc.gatech.edu
WWW URL: http://www.cc.gatech.edu/~rama
WWW URL: http://www.cc.gatech.edu/fac/Gregory.Abowd

Final Report via Fastlane November 30, 2008

This is the final report of the project. In fact, the annual report filed on August 31, 2007 represents pretty much the final report. We requested and were granted a NCE to expend the REU supplement money on aspects related to the project.

For the sake of completeness, we are repeating much of the text from the annual report filed on August 31, 2007. We have an explicitly marked addendum for the effort related to engaging REUs in the last one year at the end of this report.

1 Activities and Findings

1.1 Research and Education

The vision behind this proposal is the potential for merging the middleware needs of ubiquitous computing and grid computing. Ubiquitous computing and grid computing technologies have been maturing in parallel. The former has its beginnings in Mark Weiser's vision proposed in his 1991 Scientific American article on "The Computer for the 21^{st} Century". Grid computing has its roots in traditional high-performance computing, with initial efforts focussed on scientific and engineering applications. It was fairly straightforward to design grid middleware services for this class of applications. However, the space of applications that can benefit from the grid is expanding enormously. For example, multimedia and stream-oriented applications are becoming commonplace in the computing landscape. Quite often, end users of such applications may need support for interaction and mobility. Thus it is clear that the middleware technologies for the grid has

*HP labs, Palo Alto, CA

to evolve to account for such new demands.

Ubiquitous computing, despite its early roots stemming from Weiser's vision, has started taking off only much more recently. There are a variety of reasons for this lag, the most important of which is perhaps the lack of hardware substrates to realize the vision. With the explosion of computing and communication infrastructure in recent times, and the increasing availability of small form-factor devices there is an increase in the maturity and variety of applications. However, most of the middleware efforts in this space have been *ad hoc*, exploring only a single application or a well-defined class of applications. One of the key properties of applications belonging to the ubiquitous computing landscape is *location awareness*. A second key property is varying resource needs depending on the type of data and the sophistication of the interaction demanded by the application. Such variations may be necessitated because of user input (here the term "user" refers to human in the loop or a program) or due to system perception of the environment. Thus the middleware technology for this space has to be location-aware, dynamic, and highly adaptive. Further, as the resource demands of ubiquitous computing applications increase, it is becoming evident that the grid infrastructure would be an ideal vehicle for parceling out compute intensive parts of the application. Interestingly, the characteristics of middleware systems for supporting ubiquitous computing are not that different from those needed for supporting evolving interactive stream-oriented multimedia applications on the grid.

The research carried out in this project explored middleware technologies to support emerging ubiquitous computing applications. With growing similarities in requirements across these emerging applications and stream based, interactive applications on the grid, we target middleware research contributions that leverage existing technologies and benefit both spaces. Specifically, our top-level contribution goals are novel techniques for attribute-based naming and discovery of resources; higher-level data fusion; recovering from real and perceived failures; highly dynamic resource allocation strategies using time-varying resource loads, location data, and user demands; and detailed scalability studies of the proposed techniques in the context of ubiquitous computing applications.

This subsection summarizes the accomplishments of this project.

1.1.1 MediaBroker++ - Stream Computing on HPC

The PI (Ramachandran) working with graduate students (Dave Lillethun and David Hilley) and an undergraduate (Seth Horrigan) has developed *MB++* [5], a system that caters to the dynamic needs of applications in a distributed, pervasive computing environment that has a wide variety of devices that act as producers and consumers of stream data. The architecture encompasses several elements: The type server allows clients to dynamically inject transformation code that operates on data streams. The transformation engine executes dataflow graphs of transformations on high-performance computing resources. The stream server manages all data streams in the system and dispatches new dataflow graphs to the transformation environment. We have implemented the architecture and show performance results that demonstrate that our implementation scales well with increasing workload, commensurate with the available HPC resources. Further, we show that our implementation can exploit opportunities for parallelism in dataflow graphs, as well as efficiently sharing common subgraphs between dataflow graphs.

1.1.2 MobiGo - Seamless Mobility

The PI (Ramachandran) working with graduate student (Xiang Song) and undergraduate (Sam Young) has developed *MobiGo* [15], a system to support seamless mobility in a ubiquitous setting. Nominally, one can expect any user of modern technology to carry a handheld device such an iPAQ or cellphone and utilize resources in the environment to remain connected and enjoy continuous services while traveling from one place to the other. We present an middleware infrastructure, called MobiGo, that provides seamless mobility of services among these ubiquitous environments. We identify three different kinds of environments (spaces), self-owned, familiar, and totally-new and three axes to supporting mobility, namely, hard state, soft state, and I/O state in these spaces. MobiGo provides the architectural elements for efficiently managing these different states in the different spaces. Focusing on a specific demanding application (video service), we describe an implementation of this architecture and performance results that show that MobiGo enhances user experience for seamless mobility of video.

1.1.3 Powerline Event Detection

The PI (Abowd) with his graduate students (Shwetak Patel, et al.)) has been studying the use of the natural infrastructure in the home (such as electrical power lines) as a means of detecting the activities in the home and predicting activities. One of their papers talks about the infrastructure itself [10] and the second studies event detection using this infrastructure [9]. The latter won *best paper award* in Ubicomp 2007.

1.1.4 Software Stack for Ubiquitous Computing

The PIs (Ramachandran and Abowd) working with graduate and undergraduate students have developed *UbiqStack* [6], a taxonomy for describing the software architecture of ubiquitous computing applications. Through the lens of the UbiqStack taxonomy we survey a variety of subsystems designed to be the building blocks from which sophisticated infrastructures for ubiquitous computing are assembled. Our experience shows that many of these building blocks fit neatly into one of the five UbiqStack categories, each containing functionally-equivalent components. Effectively identifying the best-fit Lego pieces, which in turn determines the composite functionality of the resulting infrastructure, is critical. The selection process, however, is impeded by the lack of convention for labeling these classes of building blocks. The lack of clarity with respect to what ready-made subsystems are available within each class often results in naive re-implementation of ready-made components, monolithic and clumsy implementations, and implementations that impose non-standard interfaces onto the applications above. This paper describes the UbiqStack classes of subsystems and explores each in light of the experience gained over 2 years of active development of both ubiquitous computing applications and software infrastructures for their deployment.

1.1.5 Streaming Applications and the Grid

Streaming applications such as video-based surveillance, habitat monitoring, and emergency response are good candidates for executing on high-performance computing (HPC) resources, due to their high computation and communication needs. Such an application can be represented as a coarse-grain dataflow graph, each node corresponding to a stage of the pipeline of transformations that may be applied to the data as it continuously streams through. Mapping such applications to HPC resources has to be sensitive to the computation and communication needs of each stage of the pipeline to ensure QoS criteria such as latency and throughput. Due to the dynamic nature of such applications, they are ideal candidates for using ambient HPC resources made available via the grid. Since grid has evolved out of traditional high-performance computing, the tools available, especially for scheduling, tend to be more appropriate for batch-oriented applications. We have developed a scheduler, called *Streamline* [1], that takes into account dynamic nature of the grid and runs periodically to adapt scheduling decisions using (a) application requirements (per-stage computation and communication needs), (b) application constraints (such as co-location of stages on the same node), and (c) current resource (processing and bandwidth) availability. The placement generated by Streamline heuristic can be used by other services (such as task migration) in a grid environment in order to dynamically adapt the performance of a streaming application.

We have designed our scheduling heuristic over the existing grid framework, using Globus Toolkit. In our experimental study, we compare Streamline with the Optimal placement (for small dataflow graphs) and approximation algorithms using Simulated Annealing (for both small and large dataflow graphs). Most of the existing grid schedulers such as Condor, Legion and Nimrod-G focus on allocating resources for batch-oriented applications. We analyze how such existing schedulers in grid can be enhanced to support streaming applications using Condor as example. Condor is a well studied resource allocator for grid and it uses DAGMan for task graph based applications. DAGMan is designed for batch jobs with control-flow dependencies and ensures that jobs are submitted in proper order, whereas different stages of a streaming application work concurrently on a snapshot of data. Thus, we have extended Condor to meet the particular streaming requirements, resulting in a baseline scheduler called *E-Condor*.

We compared the performance of Streamline with Optimal, Simulated Annealing, and E-Condor for “kernels” of streaming applications. The results show that our heuristic performs close to Optimal and Simulated Annealing, and is better than E-Condor by nearly an order of magnitude when there is non-uniform CPU resource availability, and by a factor of four when there is non-uniform communication resource availability. We have considered two variants of Simulated Annealing algorithm with different execution times and observe that neighbor-selection and annealing schedule in a Simulated Annealing algorithm have a relatively higher impact on the performance of generated schedule for communication-intensive kernels than for computation intensive kernels. We have also conducted scalability studies and demonstrate the scalability of our heuristic for handling large-scale streaming applications. The results show that Streamline is more effective than E-Condor in handling large dataflow graphs, and performs close to Simulated Annealing algorithms, with smaller scheduling time.

While Streamline does the placement for such streaming applications well, it is clear that the application dynamics may result in the computation and communication characteristics of the application changing over time. Perhaps even the dataflow graph of the application could change over time with the addition and deletion of new stages to the pipeline. If the application characteristics that are profiled and used in the placement are considered the “typical” (for e.g. mean) values for the respective stages, then the mapping given by Streamline would result in an acceptable level of performance despite this dynamism. Nevertheless, it is important to consider the impact of the application dynamism and the consequent adaptation of the scheduling heuristic. Such an adaptive scheduling heuristic is part of our future work. We are also looking at using the placement generated by Streamline heuristic in other grid services to dynamically adapt the performance of an application.

1.1.6 GridLite

We have developed an infrastructure called *GridLite* [14] for provisioning and managing grid services in resource constrained devices. This was work done by graduate student Xiang Song while he was an intern with Co-PI Dr. Raj Kumar of HP Labs.

GridLite is an extensible framework that provides services to users on ubiquitous, resource-limited devices within a Grid infrastructure. It uses a server infrastructure for provisioning of persistent services, and smart helper services running on the “lite” devices which tap into this infrastructure. One of the goals of GridLite research is to define a Grid architecture which manages these devices such that their resource constraints are minimized by the intelligent Grid infrastructure. This is done by defining new services for managing various resources. In this work, we discuss the background, requirements and architecture of GridLite. We also demonstrate a proof of concept by implementing GridLite architecture in our test bed.

This work was also highlighted in the Global Consortium Journal in 2005 [3].

1.1.7 Support for Interoperability in Pervasive Systems

Dr. Jin Nakazawa, a visiting research scholar from Keio University, is associated with Professor Ramachandran’s research group. He has developed a universal middleware called *uMiddle* [7] that allows devices from different protocol families (UPnP, Bluetooth, etc.) to be used without any special configuration in a pervasive computing application. This past year, working with the PI (Ramachandran) and Professor Keith Edwards of Georgia Tech, and Professor Hideyuki Tokuda of Keio University, Dr. Nakazawa has designed a *A Bridging Framework for Universal Interoperability in Pervasive Systems*.

This work explores design patterns and architectural tradeoffs that arise when trying to achieve interoperability across communications middleware platforms, and describes uMiddle, a bridging framework for universal interoperability that enables seamless device interaction over diverse middleware platforms. The proliferation of middleware platforms that cater to specific devices has created isolated islands of devices with no uniform protocol for interoperability across these islands. This void makes it difficult to rapidly prototype pervasive computing applications spanning a wide variety of devices. We discuss the design space of architectural solutions that can address this void, and detail the tradeoffs that must be faced when trying to achieve cross-platform interoperability. uMiddle is a framework for achieving such interoperability, and serves as a powerful platform for creating applications that are independent of specific underlying communications platforms.

1.1.8 G^2UI

Dr. Jin Nakazawa, a visiting research scholar from Keio University, working with the PI (Ramachandran) and Professor Keith Edwards of Georgia Tech, and Professor Hideyuki Tokuda of Keio University, developed *GeoGraphical User Interface (G^2UI)*, hat enables intuitive, adaptive, and configurable multimedia exploration applications using diverse devices. While the rapidly expanding variety of multimedia devices facilitates capturing real world events through various digital media, the interface for playback of these digital media is not well integrated into the physical world. G^2UI addresses this difficulty by offering a novel way of combining physical artifacts (such as paper maps) with off-the-shelf pointing devices such as wireless mice. G^2UI facilitates storing, retrieving, and processing of multimedia digital content by the colocation of a mouse on a two-dimensional physical surface such as a geographical map. A powerful and extensible middleware infrastructure that allows the seamless interoperability of a variety of multimedia platforms and communication regimes (such as Bluetooth and UPnP) underlies this UI paradigm. The combination of the novel UI and the extensible middleware infrastructure facilitates a rich variety of geographical multimedia explorations. Details of the G^2UI paradigm, its implementation, and several application instances to illustrate the power of G^2UI form the main focus and contribution of this work.

1.1.9 Application Development

Over the life of the grant, we have developed several applications as drivers for this work. These include:

- **TVWatcher.** With the explosion of streaming content in broadcast media, there is a need for a system architecture that automates the capture, filtration, categorization, correlation, and higher level inferencing of such data from distributed sources. TVWatcher is a prototypical example of an application that demonstrates all of the above needs. This application allows user-controlled correlation of live television feed and enables a user to automatically navigate through the available channels to choose the content of interest. Symphony is an architecture for the distributed real-time media analysis and delivery which meets the system requirements for such applications. TVWatcher is built on top of the Symphony architecture, and currently uses closed-captioning information to correlate television programming. Through user studies we show that correlation engine is able to consistently pick significantly useful and relevant content. TVWatcher is an example of a stream-based, highly interactive pervasive computing application that requires sophisticated processing. A small deployment of this application may only need a cluster to perform video and audio analyses in realtime to support a few users' interactive correlation demands. To support greater scalability and applicability, we are using TVWatcher as a primary motivating application for our research into pervasive grid technologies such as a highly dynamic scheduling service.
- **EventWeb.** While the volume and diversity of multimedia permeating the world around us increases, our chances of making sense of the available information do the opposite. This environment poses a number of challenges which include achieving scalability while accessing all the available media, live and archived, inferring its context, and delivering media to all interested parties with its context attached. We envision a solution to this set of challenges in a novel system architecture. As a starting point, however, we select a previously described framework, EventWeb, suitable for annotating raw multimedia data with context meaningful to end users. We then map it onto a distributed architecture capable of correlating, analyzing, and transporting the volumes of data characteristic of the problem space. This paper first presents the requirements for our architecture, then discusses this architecture in detail, and outlines our current implementation efforts. EventWeb vision of wide-area data capture, distribution, processing and storage in realtime immediately motivates the usage of existing grid technologies for managing resources. EventWeb is another primary motivating application for our pervasive grid research, and we are federating our MediaBroker architecture and implementation to support wide-area, grid-based continuous query processing of streaming data.
- **Family Intercom.** The Family Intercom application facilitates communication between residents in a home. Our intercom implementation, built using the MediaBroker for stream discovery, distribution and processing, leverages the UPnP-based locationing system also used for the Digital Family Portrait application. It is implemented as an Intercom Control Client and many Intercom Clients spread throughout the house. Each Intercom Client accepts commands from the Intercom Control Client telling it which source its sink should listen to. As communicating users move throughout the home, the Intercom Control Client receives location updates from the locationing system and dynamically reroutes the effective audio communication by sending disconnect and connect commands to the appropriate Intercom Clients. We have implemented the application and are performing scalability studies with respect to the latency of rerouting, and with respect to an older monolithic implementation that does not use the UPnP-based locationing system. Once the MediaBroker substrate becomes federated, we believe a more large scale deployment of this teleconferencing infrastructure using fed-

erated Stampede and dynamic grid-based resource allocation and scheduling would enable far greater scalability and applicability. Furthermore, using the built-in authentication services provided by the grid toolkit would assist in securing communication channels used by the Family Intercom.

1.2 Training and Development

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.

Two undergraduates (Seth Horrigan and Ryan Juang) worked on the MediaBroker project. Seth is a co-author on a paper that appeared in RTCSA 2007. Another undergraduate (Sam Young) worked on the MobiGo project.

Two Masters students (Saurab Shah and Tushar Bansal) signed up to do special problem in Spring 2006. One student implemented an index service for use with the Globus toolkit. Another student implemented components of a pervasive computing application (such as face detector) as individual services for use in a dynamic service composition framework.

Bikash Agarwalla, a graduate student supported by this award, is doing his dissertation research on *Streaming Grid*. This research concerns supporting streaming applications typical of pervasive computing on grid infrastructure. he is in the final throes of his graduate work and is set to graduate in Fall 2007. He has accepted a position with “Ask.com”.

Xiang Song, another graduate student supported by this award, is doing his dissertation work on service virtualization for mobile personal devices (such as handhelds). He has completed his Ph.D. proposal in summer 2007 and is expected to graduate in Spring 2008.

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 NSF I/UCRC award and in relation to this award we have an industrial advisory board comprising of researchers from all the leading industries. We have two board meetings every year, 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 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. Graduate students working with the PI (Ramachandran) have year-round internships at FRB-Atlanta 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.

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.

3 Contributions

The activities of this project 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

Two PhD level students are directly supported by this award. However, due to other supplementary grants and student interest in this project, we have roughly 6 graduate students associated with the project at any point of time. Additionally, we have a good number of undergraduates who are associated with the project as well, roughly 6 at any point of time.

3.2 Student Placement

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

Bikash Agarwalla, who is completing his PhD in August/September 2006, has accepted a position at "Ask.com".

Scott Saponas, an UG alumnus from our project, is currently a graduate student at University of Washington, Seattle.

Sam Young, another UG alumnus from our project, is currently employed in Microsoft.

Seth Horrigan, another UG alumnus from our project, has been accepted for graduate studies at several leading institutions including UC-Berkeley and will be joining one of them in January 2008.

PhD student Xiang Song has spent quite a bit of time working with Dr. Raj Kumar at HP Labs in Palo Alto on grid infrastructure for "appliance computing." He spent the summer of 2007 at Samsung Electronics in Suwan, South Korea.

PhD student Dave Lillethun has been invited to spend 3 months internship at Microsoft Cambridge starting in February 2008.

4 Special Requirements

We had completed the objectives of the project as of August 31, 2007, and all the funds originally allocated to the project were also fully expended at that time, except for some remaining funds in an REU supplement to this award.

To give exposure to the research artifacts from this completed project (such as scheduler extensions to GRID computing) to undergraduates, we requested and were granted a no-cost extension of the award to support undergraduates under this REU supplement.

5 Addendum for 2007-08

As we mentioned in the beginning of this report, we had completed the goals of the project at the time of filing the annual report on August 31, 2007.

This section chronicles the engagement of undergraduates in the project via the remaining funds in the REU supplement.

I teach an UG systems course CS 2200 - Introduction to Computer Systems and Networking, typically taken by UG in their sophomore year. I use this course as a recruiting tool for drafting the cream of the crop for research exposure in my lab. I taught this in Spring 2007 and Spring 2008, and recruited several such students. I get these students to work in apprentice mode with my senior PhD students.

In this past year, we have entertained a number of UG students as researchers in our lab. They include:

- Michael S. Slaughter (worked with Nova Ahmed on RF^2ID project in Fall 2007 and Spring 2008), Man Fung Yip (an exchange from Hongkong who worked with Nova Ahmed on RF^2ID project in Spring 2008).
- Matthew Wesley McCauley, Joseph Woo, and Robert Parker (worked with Junsuk Shin on ASAP project in Summer 2008).

Incidentally, one of the graduate students supported by this grant, Xiang Song, successfully defended his thesis and graduated in Summer 2008. Presently, he is working for Microsoft Corp., in Redmond, Washington.

References

[1] 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.

[2] 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.

[3] 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.

[4] 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.

[5] David J. Lillethun, David Hilley, Seth Horrigan, and Umakishore Ramachandran. Mb++: An integrated architecture for pervasive computing and high-performance computing. In *The 13th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA 2007)*, August 2007.

[6] 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.

[7] 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.

[8] 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.

[9] S. N. Patel, T. Robertson, J.A. Kientz, M. Reynolds, and G.D. Abowd. At the flick of a switch: Detecting and classifying unique electrical events on the residential power line. In *Ubicomp 2007*, September 2007.

[10] S. N. Patel, K.N. Truong, and G.D. Abowd. Powerline positioning: A practical sub-room-level indoor location system for domestic use. In *Ubicomp 2006*, September 2006.

[11] 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.

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

[13] 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.

[14] 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.

[15] Xiang Song and Umakishore Ramachandran. Mobigo: A middleware for seamless mobility. In *The 13th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA 2007)*, August 2007.

[16] 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.