

# Advanced media-oriented systems research: ubiquitous capture, interpretation, and access

College of Computing  
Georgia Institute of Technology  
Annual Report via Fastlane July 28, 2000

**Principal Investigators:** *Umakishore Ramachandran, Mustaque Ahamad, Chris Atkeson, Ann Chervenak, and Karsten Schwan*

**Contributing Investigators:** *Gregory Abowd, Irfan Essa, Richard Fujimoto, Ken MacKenzie, Ed Omiecinski, and Sudhakar Yalamanchili*

**Other Senior Personnel:** *Calton Pu, and Ling Liu*

## 1 Activities and Findings

### 1.1 Research and Education

We are using two applications as drivers for this research. The first application is a university-wide education repository called *Intelligence Quotient (IQ)*, an outgrowth of *Classroom 2000*, a system developed at Georgia Tech for capturing in-class experiences in multiple media formats for later use. The IQ application builds on the Classroom 2000 prototype. It will have extensions for virtual classrooms, and distance education or industrial training, when we no longer want to assume that all students are co-located with one another or the instructor in space and/or time. The capabilities of IQ will include (a) support for students' interactive presence from remote locations, (b) support for personal interactive annotations to on-line course material, and (c) access to other related course material (including those in archives as well as concurrent with a particular course).

The second application is more exploratory in nature, and is called *computational spaces*. Specific instances of such applications include tele-medicine, interactive physical/sports training, distributed games and simulations, and collaborative interactions in science and engineering. The system-level and application-level software infrastructure for such applications will capture various media streams (video, audio, sonar, etc.), and/or generate such streams (e.g. scientific data), and process them for interaction, interpretation, transformation, transmission, and display without requiring explicit actions by end users.

There are two major attributes that are common to the educational repository IQ, and the computational spaces: (1) they require the capture and interpretation of live mediadata from multiple sources such as cameras, microphones, and touchscreens, and (2) they need efficient and timely access to such data from anywhere in the geographical region of interest.

Figure 1 shows the overall architecture of the proposed research. Our application- and system-level software building effort will be on top of standard operating systems (such as Unix and NT) and standard graphical and media interfaces.

#### 1.1.1 Curriculum Development

One of the pedagogical goals of the project is to involve more undergraduates in research. To this end, we have re-designed the undergraduate curriculum front-loading it with required courses in the first two years and electives in the junior and senior years. This curriculum is in place as of the Fall of 1999. As an example, on the systems side we have a sophomore level course that presents the architectural components of a computer together with the operating systems concepts for managing the resources. This method seems

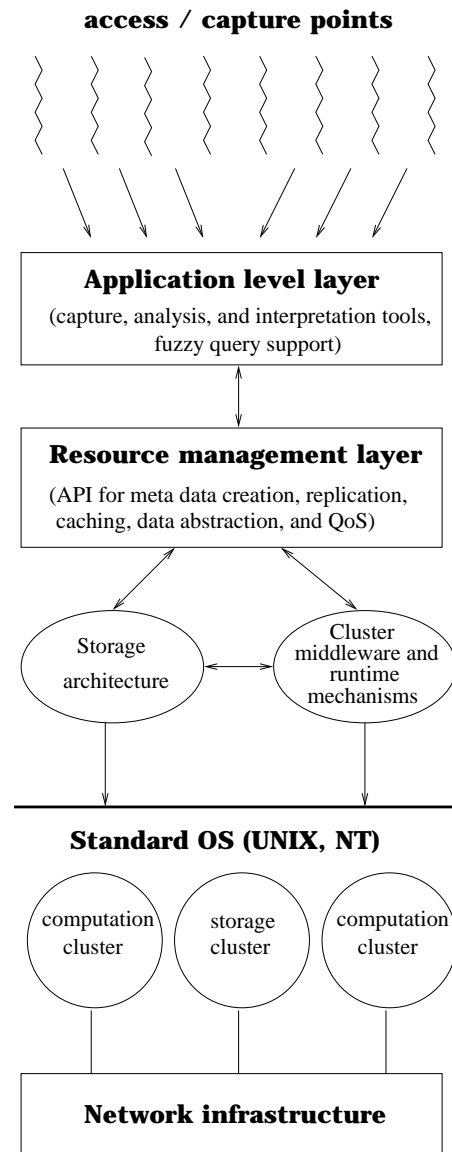


Figure 1: Overall Architecture of the Proposed Research.

very promising in revealing the synergy between hardware and software for the students. This course is followed in the junior year by two courses one in software and the other in hardware. The software course is “Operating Systems Design”, wherein we give a tour of Linux operating systems kernel internals, and have the students modify aspects of the kernel (e.g. scheduling) on a standalone processor board (the Skiff personal server designed by our industrial partners Compaq CRL). The hardware course is “Processor Design”, wherein we have students design a processor to run JVM in a couple of different styles (a conventional pipelined processor, and a special-purpose design) and compare them quantitatively. These courses should prepare the students to get involved in systems research in their senior year and accomplish significant research artifacts.

## 1.2 Major Findings

Our grant started September 15, 1999. The activities in the first year have mostly been around developing the systems laboratory and the acquisition of the high-end interactive equipment requested in the proposal. There has been progress along the research front as well on aspects of the proposed research.

We have had a significant number of research results in the scope of the RI project that span cluster computing and support for interaction. These results have often been in partnership with our industrial collaborators

such as IBM, Intel, Compaq, and Wind River Systems. The references in the publications portion of this annual report give pointers to some of these results. For a more complete list please visit the project web site at [www.cc.gatech.edu/~rama/nsf-ri/](http://www.cc.gatech.edu/~rama/nsf-ri/)

### **1.3 Training and Development**

Several undergraduate and graduate students have been brought on board for this project. They are working on different aspects of the project. Several students have been trained on the Compaq's skiff personal server which is used as one of the capture/access points for multimedia data.

### **1.4 Outreach**

The College of Computing at Georgia Tech has been running a summer intern program (funded by ONR) for promising undergraduate students from Colleges and Universities who do not have the infrastructure to support research in computing. For example, for the summer of 2000, we have participants from Spelman College, University of Puerto Rico, Florida A&M, Talladega College, Morris Brown College, and Morehouse College. The intent is to match these students with faculty and graduate students who can mentor them through an eight-week research project. This program provides the interns with a rich summer research experience, and serves as a port-hole for what awaits them if they decide to go to graduate school.

The labs mentioned above will be made available to such students to enrich their experience at Georgia Tech.

#### **1.4.1 Additional Funding**

The NSF RI award has been instrumental in attracting additional personnel funding from industry and federal and state agencies. Here is a partial list:

- K. Schwan, U. Ramachandran, K. Mackenzie, R. Fujimoto, S. Yalamanchili, D. Schimmel, V. Mooney, M. Swaminathan, V. Madisetty (PI), "The Yamacraw Embedded Systems Initiative", State of Georgia, approx. \$500,000, Sept. 1999 - Aug. 2000.
- U. Ramachandran, K. Mackenzie, S. DeWeerth, "A distributed programming infrastructure for integrating smart sensors", Georgia Tech Broadband Institute, \$50,000, July 2000 - June 2001.
- K. Schwan and M. Ahamad, "Object and Agent Technologies for High Performance Interaction", SUN Microsystems, Academic Equipment Grant, Dec. 1999, approx. \$100,000.
- K. Schwan, Calton Pu (PI), Ling Liu, Mustaq Ahamad, Yannis Smaragdakis, "Specialized Infopipes for I/O-Intensive Embedded Systems", DARPA, approx. \$400,000/year, July 2000 - Oct. 2005.
- K. Schwan and B. Plale, "Profiling the Performance of SMP Servers", Intel Corporation, Jan. 1999, approx. \$32,000.
- K. Schwan, C. Pu (PI), L. Liu, M. Ahamad, Y. Smaragdakis, "Specialized Infopipes for I/O-Intensive Embedded Systems", DARPA, approx. \$400,000/year, July 2000 - Oct. 2005.
- K. Schwan, C. Pu (PI), J. Walpole, L. Liu, and D. Steere, "Infosphere: Smart Delivery of Fresh Information", subcontract on runtime quality management, DARPA, approx. \$450,000/year, Jan. 2000 - Oct. 2002.
- U. Ramachandran, "Resource Management Issues in Stampede", COMPAQ CRL, \$20,000, Jan 1999.
- U. Ramachandran, "Ubiquitous Presence", COMPAQ CRL, \$50,000, June 1999 - June 2000.
- U. Ramachandran, "Computing on Clusters: Data Abstractions, Runtime Systems, and Resource Management", NSF, \$230,000, Sept 199 - Aug 2002.

## **2 Publications and Products**

### **2.1 Publications**

See the references at the end of this document.

### **2.2 Web Site**

Please visit the project web site at [www.cc.gatech.edu/~rama/nsf-ri/](http://www.cc.gatech.edu/~rama/nsf-ri/)

### 3 Contributions

Through the proposed research we expect to make significant contributions in the following areas:

- techniques and algorithms for machine perception of complex human behavior; tools for the ubiquitous capture and access to multimedia information; and techniques for extracting content from and supporting fuzzy queries in aware homes, in educational settings, and in perceptive intelligent kiosks,
- metadata creation to aid in inter-relating independent information streams,
- Quality of Service (QoS) architecture for adapting the application and the allocation of distributed resources to suit the dynamic needs of the applications and the availability of the computational, storage, and communication resources in the system,
- QoS-sensitive distributed middleware for consistency management of the shared computation state to increase performance and availability,
- data organization in multi-level storage hierarchies respecting the real-time capture requirements of interactive streams, and techniques for replication and prefetching to aid in ubiquitous access to captured streams,
- data abstractions and runtime mechanisms for interactive stream-oriented parallel programming on computational clusters, and hardware enhancements to device interfaces for supporting stream-oriented computations,
- runtime mechanisms (such as load balancing, scheduling, and parallelism management) in distributed clusters to support the capture, interpretation, inter-relation, and access to large numbers of simultaneous media streams, and
- design and evaluation of a distributed clustered system in the context of such emerging interactive and computationally intensive application classes.

#### 3.1 Human Resources

We have roughly 25 graduate students, 10 undergraduate students, 4 research scientists, and 1 post doctoral fellow working in areas related to this project. Of these 1 research scientist (Beth Plale), three graduate students (Vernard Martin, Durga Devi Mannaru, and Tianying Chang), and two undergraduates (Hua Yang, Shweta Humad) are from under-represented groups. With the funding available from this award we have hired a part-time research scientist (Matt Wolfe) who has been overseeing all the equipment acquisitions for this year. We have also hired an undergraduate research assistant (Matt Wolenetz) to help with overseeing the inventory of acquisitions being made for the project, and enable their use by the project participants. We plan to hire a research scientist fully dedicated for this project in the coming year.

Several students have been trained on the Compaq's skiff personal server which is used as one of the capture/access points for multimedia data. Besides, software artifacts developed as part of the project are being used in graduate and undergraduate courses thus infusing research back into the curriculum.

#### 3.2 Laboratory and Testbed Development

As part of the matching commitment to the NSF RI award, the College of Computing has made additional lab space available for housing the high-end interaction devices. The College also partly bore the cost of renovating the space to accommodate the equipment. The renovation is mostly complete. The present layout has 2 large labs abutting each other, and connected by a sliding glass door. The first lab is more traditional and hosts workstations for use by students. This lab also houses a 17-node cluster (each node is an 8-way SMP using 550 MHz Pentium III processor) which was donated by Intel. The second lab will house the high-end interaction devices such as an Immersadesk, and a video wall. The NSF RI equipment will be made available to users via these two labs. The users include faculty, research scientists, post-docs, and students (both graduate and undergraduate), from the College of Computing and elsewhere on campus. The equipment acquisition for year one is in progress and we expect the labs to be fully operational with the new acquisitions by the end of the summer.

In the proposal, we mention that the techniques we develop in this project will be deployed in several application testbeds. One such testbed being developed as part of the *aware home* project at Georgia Tech is the "Broadband Institute Residential Laboratory". This laboratory is now in place as of April 2000 for experimentation. One of our industrial partners (Compaq CRL) is developing instances of embedded computing

devices (such as the Skiff personal server <http://www.crl.research.digital.com/projects/personalserver/default.htm>). We will incorporate such devices in the application testbeds, and use the testbeds to evaluate our techniques.

### 3.3 Equipment Acquisition

The following acquisitions have been made for the first year to increase the infrastructure of the college and the institute. The acquisitions use both the NSF RI funding and the matching funds from the Georgia Institute of Technology.

- Upgraded our 16 processor Origin 2000 to a 16-processor Onyx2 with an IR3 graphics pipeline, digital video and audio inputs and outputs, and a full graphics visualization console.
- Purchased an Immersadesk R2 from Fakespace systems, along with a high-bandwidth video extender from lightwave communications to allow us to drive the desk from long distances.
- Upgraded the networking infrastructure to Cisco switches with 128 GB non-blocking backplanes to allow for gigabit or higher connectivity between the different pieces of research equipment (scheduled for early summer 2000).
- Engineered a large-format data wall, with a 2x2 array of 50" Plasma displays, embedded cameras for interaction, integrated room audio system, and a centralized control system to allow easy user access (scheduled for summer 2000).

## 4 Special Requiements

### 4.1 Plans for the Coming Year

We will continue to develop technologies for distributed access, capture, and interpretation of data from a variety of sources. Projects such as Infosphere and Stampede that are allied to the theme of the RI award help develop these technologies using the research equipment being acquired. We will also deploy many of these technologies in testbeds such as the *Aware home*, and instrumented classrooms.

## References

- [1] Nissim Harel, Vivekanand Vellanki, Ann L. Chervenak, Gregory Abowd, and Umakishore Ramachandran. Workload of a Media-Enhanced Classroom Server. *Proceedings of IEEE 2nd Annual Workshop on Workload Characterization*, October 1999.
- [2] K. Knobe, J. M. Rehg, A. Chauhan, R. S. Nikhil, and U. Ramachandran. Dynamic Task and Data Parallelism Using Space-Time Memory. In *SC99*, December 1999.
- [3] R. S. Nikhil and U. Ramachandran. Garbage Collection of Timestamped Data in Stampede. in *Principles of Distributed Computing (PODC 2000)*, July 2000.
- [4] J. M. Rehg, K. Knobe, U. Ramachandran, and R. S. Nikhil. Integrated task and data parallel support for dynamic applications. *Journal of Scientific Programming*, pages 289–302, 1999.
- [5] Vivekanand Vellanki and Ann L. Chervenak. A Cost-Benefit Scheme for High-Performance Predictive Prefetching. *Proceedings of Supercomputing '99 Conference*, November 1999.
- [6] G. Brostow and I. Essa. Motion-based video decompositing. In *Proceedings of IEEE International Conference on Computer Vision 1999 (ICCV'99)*, pages 8–13, Corfu, Greece, 1999. IEEE Computer Society.
- [7] A. Chervenak, V. Vellanki, I. Yanasak, N. Harel, R. Rodenstein, G. Abowd, J. Brotherton, and K. Ramachandran. A scalable workload model of media-enhanced classrooms. Technical Report GIT-CC-99-06, College of Computing, Georgia Institute of Technology, 1999.
- [8] Kidd. C., G. Abowd, C. Atkeson, I. Essa, B. MacIntyre, E. Mynatt, T. Starner. The Aware Home: A Living Laboratory for Ubiquitous Computing Research In *Proceedings of Second International Workshop on Cooperative Buildings 1999*, Editors, Streit, J. Siegel, V. Hartkopf, S. Konomi, Pittsburgh. LNCS 1670. Springer: Heidelberg, Pittsburgh, PA, October 1999.

- [9] D. Moore, I. Essa, and M. Hayes. Exploiting human actions and object context for recognition tasks. In *Proceedings of IEEE International Conference on Computer Vision 1999 (ICCV'99)*, pages 80–86, Corfu, Greece, 1999. IEEE Computer Society.
- [10] F. Bustamante, G. Eisenhauer, K. Schwan, and P. Widener. Efficient Wire Formats for High Performance Computing. To appear in *ACM Supercomputing 2000*, (SC 2000), Nov. 2000.
- [11] R. Krishnamurthy, K. Schwan, and M. Rosu. A Network Co-Processor-Based Approach to Scalable Media Streaming in Servers. To appear in *International Conference on Parallel Processing (ICPP)*, August 2000.
- [12] M. Rosu and K. Schwan. Support for Recoverable Memory in the Distributed Virtual Communication Machine. To appear in *International Parallel and Distributed Processing Symposium (IPDPS)*, May 2000.
- [13] B. Plale and K. Schwan. dQUOB: Efficient Queries for Reducing End-to-End Latency in Large Data Streams. To appear in *High Performance Distributed Computing (HPDC)*, August 2000.
- [14] G. Eisenhauer, F. Bustamante, and K. Schwan. Event Services for High Performance Computing. To appear in *High Performance Distributed Computing (HPDC)*, August 2000.
- [15] B. Plale, G. Eisenhauer, L. K. Daley, P. Widener, and K. Schwan. Fast Heterogeneous Binary Data Interchange for Event-based Monitoring. To appear in *ISCA 13th International Conference on Parallel and Distributed Computing Systems (PDCS 2000)*, August 2000.
- [16] C. Isert and K. Schwan. ACDS: Adapting Computational Data Streams for High Performance. *International Parallel and Distributed Processing Symposium (IPDPS)*, May 2000.
- [17] G. Eisenhauer, F. Bustamante, and K. Schwan. A Middleware Toolkit for Client-Initiated Service Specialization. *Principles of Distributed Computing Middleware Symposium (PODC 2000)*, July 2000.
- [18] D. Ivan-Rosu and K. Schwan. FARACost: An Adaptation Cost Model Aware of Pending Constraints. *Real-time Systems Symposium (RTSS)*, IEEE, Dec. 1999.
- [19] L. Liu, C. Pu, K. Schwan, J. Walpole. InfoFilter: Supporting Quality of Service for Fresh Information Delivery. *New Generation Computing Journal*, Ohmsha, Ltd. and Springer-Verlag, Special issue on Advanced Multimedia Content Processing, Vol.18, No.4, to appear August, 2000.