Advanced Media-oriented Systems Research at Georgia Tech: A Retrospective

Umakishore Ramachandran, Karsten Schwan, Phillip Hutto, Matthew Wolf

College of Computing, Georgia Institute of Technology, Atlanta, Georgia
{rama, schwan, pwh, mwolf}@cc.gatech.edu
http://www.cc.gatech.edu/~rama/nsf-ri

Abstract

This paper reviews the six year history (September 1999 through August 2005 with one year no-cost extension) of advanced media-oriented systems research in the College of Computing at the Georgia Institute of Technology with respect to the execution of NSF CISE funded grant EIA-9972872 "Advanced Media-oriented Systems Research: Ubiquitous Capture, Interpretation, and Access". We describe the original goals of the project and discuss the evolution and adaptation of the research agenda in response to changing personnel and interests, research experiences, and emerging technologies.

We offer several "lessons learned" from outfitting and working with smart spaces and pervasive computing infrastructure and highlight, as a general theme, the inevitable tension that arises from demands for "production access" to successful experimental research facilities. We conclude with some general comments about the fortuitous research outcomes that have resulted from the varied resources, tensions, and pressures that characterize a large, multi-year equipment grant involving a large number of researchers with diverse research agendas. In particular, we believe we have gained considerable insight into the integration of "large and small" computing elements in service of pervasive computing applications.

1 Evolution of Research Goals

Our original grant proposal[1], authored in early 1999, described our objectives as follows:

"Using two large-scale research applications—a distributed education repository, and perceptual computational spaces for multimedia-based collaboration—as drivers, we propose to carry out extensive systems research and integration to support ubiquitous access, capture, and interpretation of a variety of multimedia streams."

We anticipated increasing demand for rich access and manipulation of complex data and media as sensors and actuators and computational resources were ubiquitously deployed and high-speed Internet access and wireless became commonplace. We believed that the demand for such access would give rise to a new class of interactive distributed applications that would require sophisticated real-time coordination (fusion, correlation, sampling, transformation) of multiple media streams and other existing data sources. We aimed to engage in end-to-end analysis of the system infrastructure required in such a world, driven by experience with real-world applications.

At the time of this writing, the British government is currently engaged in a massive effort to review hundreds of hours of video and other forms of media and data in order to bring to justice those involved in the subway and bus bombings of July 8th, 2005. This is just one rather grave example demonstrating the importance of the application domain we have envisioned.

While one of our driver applications (the distributed education repository) was dropped early on due to the vagaries of campus politics, and while we have undergone some personnel changes (two of the original PIs, Chris Atkeson and Ann Chervenak, are no longer at Georgia Tech), our predictions about the increasing importance of this application domain seem to be on target.

We have remained true to the intent of the original proposal, undertaking research of many infrastructure components (e.g. DFuse, Media Broker, Energy-Aware Traffic Shaping and Transcoding, Stream Scheduling, Agile Store, Differential Data Protection) involving many different driver applications (e.g. Aware Home, Smart Spaces, High-Performance Computing Program Steering, Event Web, TV Watcher, Vehicle-to-Vehicle Networks), and involving many application domain collaborators (e.g. Abowd, Essa, Fujimoto, Rehg, Starner and Jain).

2 Lessons Learned Outfitting Smart Spaces

Room 207 (the Systems Studio) in the College of Computing Building (CCB) was outfitted at grant inception as a "smart conference room" to support project research efforts and included audio and video support with ceiling microphones and speakers, overhead projectors, a video wall, cameras, an Immersadesk for 3-d visualization, a variety of desktop systems and a rich interconnect and processing infrastructure with touch control panel. This facility was connected via high-speed networking and wireless to various servers, a large disk array and various computational clusters. A touch-screen based video-kiosk was installed in an adjacent hallway.

Subsequently, a variety of sensor and actuator devices were installed or housed in the Systems Studio

[1] "Advanced Media-oriented Systems Research: Executive Summary", available at http://www.cc.gatech.edu/~rama/nsf-ri

and adjoining areas to form the Sensor Lab, including Berkeley Motes, handhelds and laptops, IR/RF location tracking system, temperature sensors, a small mobile robot, VCRs, RFID capabilities, cameras, microphones, speakers, etc. Similar devices were also installed in the Aware Home on campus. In the summer of 2004 the facilities in the System Studio were upgraded and enhanced and a smaller conference room (CCB 256) was similarly outfitted.

2.1 Theme: Experimental or Production Facility?

A recurring theme that effected many aspects of our research efforts involved disagreement over the nature of our equipment facilities. While our smart spaces were originally intended as experimental testbeds, many forces conspired to convert them to production facilities. We believe this is a natural tension that will arise with long-lived, equipment-based grants that seek to involve diverse personnel.

Experimental facilities are, by nature, flexible and frequently undergo transformation and adaptation. They are often re-configured to host new uses and to try out new capabilities. Such re-configuration takes time and often difficulties are encountered, temporarily disabling the space. Moreover, such spaces often include new equipment that requires exploration to fully understand and utilize or they are controlled by experimental software that may be limited in capabilities, fragile, or difficult to use. In short, experimental facilities are often "broken" as a consequence of their nature. Production facilities, on the other hand, must be accessible, friendly, easy-to-use, reliable and highly-available.

Smart spaces, when successfully implemented, are appealing locations for meetings, work groups, and other events. Success in providing useful and desirable facilities and attracting "real-world users" limits the flexibility of the space for re-configuration and experimentation. Keeping the environment flexible and reconfiguring frequently to test new applications and software discourages real-world users who find the system not "setup" for their needs.

This tension plays out at many levels. It arises when determining what types of equipment to purchase (should we get a bleeding edge new device or something more "useable"?) and when determining access controls for the systems and space (should anyone be able to use the system or just authorized researchers?).

Finally, the inevitable pressures of space utilization affected our efforts. While the System Studio was originally designed as a dedicated facility, space constraints in the College of Computing required us to add general use desktop systems, complicating things when we needed to run experiments or hold Access Grid teleconferences.

2.2 Lessons Learned

We offer several suggestions for those working with similar smart spaces in the future:

1. plan ahead, get "buy in", put it in writing;
2. experimental wireless research is difficult;
3. use professional services where appropriate;
4. knowledge transfer is difficult;
5. avoid the urge to buy "cool stuff";
6. plan on upgrading for long-lived facilities

Get agreement over intended uses of the space with all relevant stakeholders and have them sign a written agreement of intent. This will come in useful down the road with changes in administration personnel. (We are happy to note that our administration has been extremely supportive and came through with all that was promised as part of institutional cost sharing.) Experimental wireless research is increasingly difficult to do with most universities (including Tech) blanketed by "standard" wireless capabilities. Some things (such as audio systems) are best done by paid professionals. Computer scientists are capable of these projects but it is best to utilize commodity goods and services when appropriate. We found that over the six year period, we would work actively with a set of equipment for awhile and then move on to other research topics. Later, new personnel would be interested in using the older equipment and have to go through a time-consuming process of re-learning how to use the equipment. Avoid the urge to buy sophisticated new technologies with the intent of understanding and integrating them later. It won't happen. Finally, if your smart space will be in use for several years, plan on upgrading the facility. We mitigated this issue somewhat by staging our purchases over the lifetime of the grant but upgrade was still required for some components to stay current.

3 Synergistic Research Outcomes

Our emphasis on application-driven research has significantly enhanced interaction among systems and application-domain researchers at Georgia Tech but these interactions were most vigorous when funded by related grants. A wide range of research results have been achieved under this grant. We believe that the synergy of equipment and researchers has helped us clarify the integration of "big and small" computational components in the service of pervasive applications with computationally intensive demands (e.g. video analysis).