# Development of a Large-Scale Ubiquitous Computing Interface

## 1 Summary

We propose the development of a unique experimental facility for the exploration of large-scale ubiquitous computing interfaces and applications that are aware of the context of their use and can capture salient memories of collaborative experiences. Our research goal is to explore new paradigms for human-computer interaction through a prototype next generation ubiquitous computing interface. Our engineering goal is to build a computational and sensing infrastructure so pervasive that it can do for supporting and capturing collaborative activities what the worldwide computer network has done for e-mail and on-line communities. Like the telephone system or the road network, this interface must be available throughout a user's work space to fully reveal how it will be used and the scientific issues raised. Current ubiquitous computing experiments are limited to the scale of a single room or just provide a minimal service (e.g., location) across a larger area. We propose developing a rich ubiquitous computing infrastructure for the distributed space occupied by the College of Computing at Georgia Tech. We will add sensing capabilities to the environment, including microphones, cameras, and position measurement systems. We will explore a variety of display elements, including speakers and handheld, wearable, and wall-mounted displays. Furthermore, we will install computational resources sufficient to process and generate all of the information that will be sensed in this environment.

Scale is not the only hard issue when dealing with ubiquitous computing interfaces. Applications supported through ubiquitous technology must provide natural, multi-modal input and output. Users may speak, gesture, or manipulate in the context of the task they are performing, and the ubiquitous interface must make sense of that input. The interface needs to be as accessible as possible so as to reduce the cost of use and support novices. These applications must also be seen by the users as providing a value-added service, or they will not be used in everyday tasks no matter how accessible the interface.

Projects that would use this experimental facility include: studies of systems that capture what is happening in an environment and make that information available to users (focusing on capture systems to enhance education); explorations of systems issues in delivering multimedia services in this type of ubiquitous computing environment; studies of how computers can be made aware of what is happening in an environment (computational perception); studies of software engineering issues for ubiquitous computing; and explorations of designs for new interfaces and interface paradigms for ubiquitous computing environments. Basic scientific questions include:

- What are the most appropriate modes of human-computer interaction in this context?
- What algorithms can support the machine perception and reasoning underlying this type of computing interface?
- How can we most effectively manage limited computational resources to provide desired services in this context?
- To what extent can the interface be automatically customized for each user using machine learning?

The College of Computing at Georgia Tech is uniquely positioned to undertake the development and investigation of a large-scale ubiquitous interface. We have a great deal of experience with ubiquitous interfaces through work done by the PIs in the Future Computing Environments research group at Georgia Tech. The Classroom 2000 project emphasizes making a room intelligent enough to capture the relevant information from a university lecture [3, 10, 4]. The Cyberguide [35, 36, 5] and CyberDesk [70, 16] projects explore ways to make the services in the system apparent to the user when and where appropriate by making the services themselves context-aware. We have also defined the topic of dual augmentation of physical and virtual worlds and demonstrated this concept through the Domisilica project [38, 40, 41]. Our systems group provides expertise in cluster parallel systems, I/O systems, large-scale storage systems and media servers. The Graphics, Visualization, and Usability Center at Georgia Tech provides strengths and expertise in user interface technologies. We will continue our collaboration with Georgia Tech's excellent DSP group. We have assembled state of the art facilities that will enable and complement the proposed instrumentation of this proposal, including the listed matching funds and an existing large-scale equipment grant from Intel Corporation.

The instrumentation from this grant will allow a large population of talented people to live and invent a paradigm shift in human-computer interaction.

# Program Description

## 2 Research Activities

### 2.1 Introduction

There is growing interest in using computing technologies to build systems that support our daily lives more effectively [26]. Examples of such systems are smart offices, classrooms, and homes that allow the computer to monitor what is happening in the environment, to figure out how to most effectively assist in the tasks being performed, and to have sufficient output devices to assist effectively. For example, a classroom might help students take more effective notes by capturing much of the details of the rich lecture experience, freeing the student to pay more attention to the content of the class and simply annotating or indexing into the record of the experience for later review. A meeting room could record the events and discussion in a meeting to allow others who did not attend the meeting to efficiently find out what happened. An office could remind its user of relevant information, or signal urgent messages. A machine room could provide documentation for operators or visiting repair technicians. A home could babysit children or allow elderly to remain self-sufficient. A building could supplement the capabilities of physically or mentally challenged individuals through visual assistance, movement assistance such as opening and closing doors, and providing and receiving information in alternate forms such as lip-reading and sign-language recognition. These systems can be built into the environment as well as being handheld or even worn as part of our clothing.

Research and development efforts for building such intelligent and interactive human-centric systems that support and augment our daily lives rely on the concepts of ubiquitous and aware computing. The defining characteristic of ubiquitous computing [67, 68] is the attempt to break away from the traditional desktop interaction paradigm (screen, mouse and keyboard); instead, computational power is embedded into the surrounding environment and is available wherever and whenever users need it. In contrast with traditional interaction where the user is forced to go to the computer interface, the ubiquitous interface itself takes on the responsibility of locating and serving the user.

The challenge of ubiquitous computing not only involves distributing the computation and networking capabilities, but also includes providing a natural interface to the user. As computation becomes more pervasive in the user's environment, the intent is to make the actual interface less obtrusive, to the point where the physical interface becomes transparent to the user.

Aware computing aims to surmise information about the user and the environment that surrounds the user. Such awareness can be achieved by incorporating perceptual abilities into the environment. This form of computational perception can be used to locate users, identify them, determine their focus of attention, and attempt to ascertain what they are doing and what their intentions are.

We are interested in combining ubiquitous and aware computing to achieve computational augmentation of our everyday activities. This coupling can be achieved by instrumenting the environment with sensors, output devices, networking, and computational power. Such instrumentation can be used for capturing and processing audio, video and other sensory data, and controlling the input, output and information flow in an environment. Present commercial and academic developments in computational hardware, input/output devices, and sensor technologies are pushing application builders to more seriously consider instrumented environments as part of the computational infrastructure surrounding a user.

Exploring new paradigms for human-computer interaction through ubiquitous computing requires full-scale prototypes. This has been a serious problem for research on ubiquitous computing interfaces, which require large areas to be instrumented. Current state of the art is to instrument a single room or a dedicated area in an office environment. We propose developing a ubiquitous computer interface for a larger area, the entire space occupied by the College of Computing at Georgia Tech. The sensing capabilities of the environment will make it easy for groups of users to return to a captured memory of some formal or informal collaborative experience. The environment will also be aware of the users and will attempt to provide services that meet their intentions. Like the telephone system or the road network, the services provided in this prototype interface must be available throughout a user's workspace in order for us to fully reveal how it will be used in the future and to identify the scientific issues raised.

### 2.2 How Would the System be Used?

The system provides several types of assistance for users:

- **Access to information:** "Show me my class notes from last term in which I learned about matrix inversion." The system can assist users in rapidly accessing material during a meeting, a class, or a conversation.
- **Communication and collaboration support:** Users should be able to see if someone is available for a virtual 'chat' with minimal effort. The system should monitor whether someone is interruptible at any given moment, and be able to provide information as to the location and status of other users, if they so desire. The system should provide subtle cues as to the status of incoming communications, and also monitor the urgency of messages or attempts to establish communication. We hope to provide a much more lightweight and flexible form of remote conferencing than the current use of special equipment in prepared locations.
- **Capturing everyday experiences** Another type of assistance is capturing what happens in the environment, integrating that information, and providing access to it later. Our initial attempt at capture is the Classroom 2000 system, described in more detail in what follows. We will extend the Classroom 2000 system to capture large and small collocated group meetings that occur everyday in the College of Computing and elsewhere, such as classes, meetings, and hallway conversations. We will also look to capture remote meetings such as video-conferences, and asynchronous activities, such as class project histories.
- **Environmental awareness:** One particular type of environmental augmentation is to provide subtle cues, or peripheral awareness, that would allow one to know what activity is underway in another space. For example, a colleague might want to monitor activity around the coffee cart to determine when would be a good time to join other colleagues for a casual conversation. We are familiar with computers using sensors to serve as burglar alarms. More intelligent systems could detect a pot boiling over on the stove or a person in need of medical assistance.
- **Automatic receptionist and tour guide:** The system could identify visitors, ask them who they wanted to see, and guide them. We have used early prototypes of our ubiquitous interface to provide tours and guidance to visitors during our open houses or visitor days. In addition, universities have a high turnover rate. Every fall we replace a quarter of our undergraduate population. Our system could help guide these newcomers.

Our goal is to support these types of tasks in offices, classrooms, common spaces, and hallways of the Georgia Tech College of Computing. As part of this goal, the interaction with the computer must appear coherent to the user, even though the user is moving from location to location, the interaction is interrupted by other activities, and there are many other users interacting with the system at the same time in the near vicinity. The College of Computing is currently spread over space in four buildings, with the furthest building a mile away. How can our ubiquitous interface help combine physically distributed spaces so that communities (formal and informal, active and passive) still thrive?

## 2.3 Research Projects That Would Use The System

We list here some of the research projects that would utilize and in fact thrive, in this experimental facility. Although these projects are listed separately, we can't emphasize enough that one of the most important functions of an experimental facility such as this one is to help motivate researchers to integrate their research, especially in an environment that is used on a daily basis.

### 2.3.1 Capture, Integration, and Access

We have been building prototypes of ubiquitous computing applications for over three years and have seen a number of common themes emerge across various applications. One of these concerns the development of automated tools to capture live experiences for the purpose of providing access to them later on. This common feature of ubiquitous computing is the subject of the NSF Career Award to Gregory Abowd, *Investigating Research Issues in Ubiquitous Computing: The Capture, Integration and Access Problem,* (IRI-9703384, Total award: $438,376 (including matching funds), 7/1/1997 – 6/31/2001). We are focused on developing *automated* techniques to capture the different multimedia streams that define a rich interactive experience. Once captured, these streams can be related, or integrated, so that information from one stream can cross-index into another stream. Later on, the relationship between streams can be revisited through universally accessible interfaces. As with most other ubiquitous computing research, our initial forays involved restricted applications.

**Capturing experiences in a classroom:** The university classroom is a restricted and useful environment to experiment with capture, integration, and access. The Classroom 2000 project (funded by the above mentioned NSF Career grant together with equipment grants from the Georgia Tech Office of Information Technology, Proxima Corporation, NEC, and Smart Technologies) has as its goal to help students and teachers function more effectively in the classroom by augmenting a single room with the capabilities to automatically take notes on behalf of the student [3, 10, 4]. The ubiquitous computing interface of the classroom includes large-scale electronic whiteboards, portable electronic student notebooks, microphones, cameras, networking, computation and mass storage. Its purpose is to allow a lecture to proceed as normal, with the end result being an automatically generated set of notes that allows the student to relive the whole lecture or more efficiently browse and navigate to points of interest. By viewing the classroom experience as a multimedia authoring session, we simplify content generation for educational software.

Over the past two years, Classroom 2000 has been used as a regular component of 25 different undergraduate and graduate courses at Georgia Tech. Section 3.2 discusses the special purpose classroom we built for experimentation. Section 4 discusses other universities and companies using our Classroom 2000 technology.

**Meetings:** Another well-defined environment is the meeting room. With sponsorship from DARPA (part of the EDCS program), the ARL Federated Laboratory on Advanced Displays, and BellSouth, we are investigating the development of a suite of applications that demonstrate the ability to automatically capture, integrate, and access strategic and technical rationale

**Mobile Exploration:** A less well defined space for ubiquitous computing is the personal space that surrounds a user. With sponsorship from BellSouth Intelliventures, we have developed a handheld mobile tour guide, Cyberguide, to assist visitors to the College of Computing and GVU Center [35, 36, 5]. A simple indoor positioning system allows the system to know where the user is located and can help to make educated guesses to inform the user about the exhibit they are nearest. After an hour or two visiting, our guests leave with an automatically generated diary of the exhibits visited in the order they were visited, annotated with comments along the way.

These examples show the general utility of a capture, integration, and access application. What we propose to do with this instrumentation grant is to provide the ability for all classrooms, all meetings (formal and impromptu) and all social interactions to be recorded, under control of the users with little to no effort. Our aim is for it to be as natural as picking up a piece of scrap paper to jot down salient notes from an important or interesting conversation.

Through our experience, we have come to understand some of the harder research issues in the general area. One is how to determine the appropriate level of granularity for multiple stream integration. In Classroom 2000, we define a simple continuum from the finest, pixel-level integration to the coarsest, lecture-level integration, as you would have if you just tape-recorded the lecture yourself. Neither extreme is very useful, but in between there are a number of useful alternatives. The best interface during the access phase would be dynamic enough to support multiple integration levels at the user's request. Having captured many different streams, it is a challenge to know how to present the experience in a way that conveys the most information and affords rapid means of search. This leads to our next research project.

#### 2.3.2 Information Management

A very important research issue raised by this large collection of captured experiences is how to augment our human ability to segment, structure, index, link, and navigate captured material. The NSF grant proposal *Automated Understanding of Captured Experience* to the Experimental Software Systems Program by the PIs is meant to examine this problem. One approach we are using is to apply statistical measures for natural language understanding [58, 7, 12, 25, 32, 49, 15, 33, 11, 59]. Statistical approaches are usually based on matching word frequencies (vocabularies) in written or spoken text. The key insight in many statistical approaches to language is that "concepts" can be represented by vocabularies. Using vocabularies enables concepts to emerge from the data, and allows us to avoid the tedious work of categorization by hand. Pieces of text which have similar vocabularies are probably about related topics, and are probably relevant to each other.

This statistical approach can: answer more general queries about captured information in a ubiquitous computing system, including trying to identify the concept the user is searching for rather than the word; automatically segment and structure captured material; automatically link captured material to other material such as textbooks, briefing materials, or web pages; and automatically find keywords or make summaries or abstracts, to aid human navigation of large amounts of material.

#### 2.3.3 Computational Perception and Aware Computing

Statistical approaches may prove adequate for facilitating search across some number of captured streams, but often critical information from the actual experience is lost, making it impossible to understand what was captured ("Which of the symbols in this equation was the professor referring to when she said 'take this term and multiply it by that term to complete the proof.’ ”?) Experimental facilities designed to capture experiences from the environment are ideal for exploration into what are the human activities in an environment, how can they be perceived, and how those perceptions can give us better information about what is actually going on. This is the subject of the NSF grant proposal *Computational Perception of Human Actions* to the CAREER program by Irfan Essa. This proposal explores new methods for machine perception, understanding, and interpretation of human activity in a scene or an environment. With a new capability of awareness, the ubiquitous interface can respond to human input where and when it is appropriate (“This is a new PhD student in the College of Computing and he looks confused. He should be directed to the front office to meet with the PhD Coordinator, his account should be set up, and campus maps should be sent to him.”)

Our interests are in understanding the activity and the actions of the people in the environment (“what is happening?”). This requires the development of algorithms for recognizing who is in the environment, and what they are doing. Users can be identified by face recognition and/or speaker identification (“who is in the space?”). After identifying the user(s), we pursue automatic interpretation of gaze and body orientation, speech attributes, facial expressions, gesture, and body movements, from the information contained in the video and audio streams. We expect the combined information from these visual and aural modalities will allow understanding, coding, and re-synthesis of the communication signals in many situations where information from only one of the modalities would not be sufficient for understanding. This level of information would be very useful in determining the focus of attention of the user and/or the audience in an environment such as a classroom. This type of work will build on our previous work on machine perception for intelligent interaction [21, 20, 22, 19, 9, 23, 50, 51].

As a test-bed for exploring this approach we are using interactive spaces, including an instrumented classroom, equipped with cameras, microphones, and recording facilities for off-line analysis of the audio and video signals. The proposed experimental facility would provide us with the ability to explore an automatic interpretation of people’s actions in environments of different types and would give us an unprecedented opportunity to evaluate the impact on real users.

### 2.3.4 Large-Scale Media Servers

Once we are able to capture all relevant information in an environment and provide efficient means to forage through those experiences, we will have created a great demand for capture, integration, and access applications. But will we be able to deploy complex distributed systems that provide the kind of performance users want? It is one thing to equip a single classroom with tremendous capabilities to capture and report effectively on lectures. In such a controlled and limited-size environment, we do not have to deal with the complexities of scale that are the bread and butter of distributed computing. The NSF proposal *A Large-Scale Media Server for Collaborative Applications* to the Experimental Software Systems Program by Abowd, Atkeson, Chervenak and Ramachandran has the goal of providing the computational infrastructure for creating and managing a complex collection of multimedia objects. If we look just at the logical extensions of Classroom 2000 across the entire College of Computing, we see a system generating content continuously, with rich and highly dynamic accesses to the server from end-user clients. The instrumentation in this proposal will serve as a realistic online test-bed for an extensive experimental program exploring performance, bottlenecks, utility, and scalability. The impact of developing the next generation of media servers that handle the creation and playback of large numbers of small interrelated objects would be a much greater penetration of media services in our daily lives.

In designing an Application Program Interface (API) for the media server, a unified approach will be taken to the creation and playback of multimedia objects, annotations, data streams, other forms of data, computational processes that act on the stored objects or produce objects or data streams, and other encapsulated computational processes. User-defined annotations to the objects stored in the repository can be of any type, and the API will present a uniform approach to labeling links to relate objects. The design will assume a highly dynamic environment, wherein new content can be continuously added and existing content can be modified or annotated. A wide variety of media types and formats, and a large number of multimedia objects from a corresponding variety of sources will be supported by the media server. The design of and the implementation choices for the media server will allow for a large number of clients to be handled simultaneously, and for a variety of interaction patterns between the clients and the server. A range of system technologies pertaining to state sharing and resource management for realizing a cluster parallel implementation of the media server will be designed and implemented. The design will allow for loads that are highly clustered in time and across objects. To the maximum extent possible, metadata and links between objects will be automatically created. The system design will accommodate loads that are highly clustered in time and across objects. The system will support more abstract queries and support complex metadata such as semantic networks. A single query (in any media) should generate results in multiple media.

There are a number of important issues we can examine for the media server. They can be broken down into questions concerning functional issues on acquiring and accessing multimedia data objects as well as engineering issues for separating application semantics from the core acquisition and access tasks of the media server. For example, when and where does processing such as filtering and speech recognition occur? Is processing applied at the sensor, at a server dedicated to that sensor, or is raw data communicated to a server associated with the user generating the data? How is such processing updated or redone when new information becomes available? As the user moves to new locations, does the processing chain change? How is the data organized for later access? How will the system be scaled up to handle more users and more input and output devices? How can we automatically configure the system? Can we define a uniform set of API calls to handle these issues? These and other similar questions will be the focus of systems faculty in the College of Computing as they integrate their work with this live ubiquitous computing system.

#### 2.3.5 Software Engineering Issues for Ubiquitous Computing

The real point of ubiquitous computing, as Mark Weiser observes, are the applications that are built to augment our everyday tasks [68]. Much of our research has focused on rapid prototyping of ubiquitous computing applications to support everyday activities. Because we have forced ourselves to use our own systems on an everyday basis, we have been faced with significant software engineering issues. We discuss three of these issues here that we are working to solve in order to realize truly ubiquitous software services.

We have identified three general functional characteristics of a number of ubiquitous computing applications —capture, integration, and access; context-aware computing; and dual augmentation [38, 40]. To support rapid prototyping, we need to provide toolkit-level support for each of these functional themes. Through matching funding from the Corporation for National Research Initiatives (CNRI), we are developing a general capture, integration, and access framework to support a wide variety of applications. Motorola is funding work on the development of a general context inferencing engine that can be factored out of the CyberDesk prototype and used in more mobile scenarios. Intel's funding of the Domisilica project for dual augmentation in the home has lead to the creation of some general mechanisms for adding a Java-based graphical interface to a traditionally text-based MUD/MOO virtual model of physical spaces.

Interoperability is a common research theme for the engineering of complex distributed systems. Solutions to this general problem can be found in networking and systems-layer research. Programming in the large has benefited from dynamic integration to allow for flexible negotiation between system components at run-time. Our CyberDesk project shows the importance and usefulness of dynamic integration of computational services [70, 16]. Dynamic integration is beneficial to the programmer because there is no need to know in advance all possible contexts of use when a service is designed. It is beneficial to the user because it allows for continuous switching and selection of different computational services. As we move to a larger ubiquitous computing testbed, the variety of different services that students and faculty will want at their fingertips will require more attention to facilitation of dynamic integration.

The instrumentation we describe in this proposal consists of interactive devices with radically different input and output characteristics. Yet, it is highly likely that users will want some number of similar services available on all of these devices (e.g., electronic mail). A challenge for ubiquitous computing is to develop applications that can automatically adjust, or scale, to these very different devices. This application interoperability problem is referred to as the scalable interfaces problem [1]. Yet another consequence of interoperability requirements is that a working ubiquitous computing system consists of a heterogeneous collection of devices with numerous failure modes. Designing a system to be robust under the inevitable threat of failures is a challenge.

#### 2.3.6 Evaluation and HCI Issues In Ubiquitous Interface Design

One of our prime motivations for proposing a system that provides ubiquitous and aware capture services is to help foster a paradigm shift in the way we as humans perceive computation. It is only through careful study of systems in everyday use that we will be able to understand how we relate to that technology. We are very fortunate at Georgia Tech to have access to researchers in the GVU Center and the EduTech Institute who can help in guiding both quantitative and qualitative studies in HCI. However, there is no body of HCI literature that suggests the correct way to evaluate ubiquitous computing. Our own experience in the Classroom 2000 project has been to employ education researchers to attend our classes and interview students and teachers.

Our philosophy of living the futures that we invent has also provided invaluable feedback into our own system designs. We have taught classes in Classroom 2000 on how to evaluate systems. At the end of a 10-week quarter, the students are asked to turn around and evaluate the technology used to assist them during the class. The students learn the benefits of longitudinal studies and we reap the benefits of a large number of well-considered evaluations on not only the interfaces provided for capture and access in Classroom 2000, but also on deeper issues having to do with how Classroom 2000 supports or does not support the learning objectives or teaching styles of a particular instructor.

One common desire for ubiquitous computing researchers is to provide transparent interfaces to the users, ones that are ever-present when we need them yet never obtrusive or invading. How can a designer build such transparency in? Just as HCI researchers have developed easy to learn and economical usability techniques to uncover the learnability errors in an interface, we must also look for ways to guide a designer toward more transparent ubiquitous interfaces. One heuristic that we have used with success is to constantly ask whether there is information in the system somewhere that can be used to remove a distracting administrative task. Writing on an electronic whiteboard should be as simple as picking up a pen and writing. The board is in a room, the room has a schedule for classes, so it should be possible to determine who is writing on the board and why.

This idea of invisibility becomes even more important in less-structured interactions. A classroom schedule is usually well-known, so it is not that difficult to see how to automate most if not all of the administrative tasks associated with capture. In a less formal setting, such as an impromptu meeting of students in a social space, it is far more difficult to support lightweight and quick interactions. It is even more difficult to detect when it is appropriate and useful to capture information from these impromptu encounters. We want a system to be able to detect when a group of students congregate, recognize who they are, and why they might be meeting. If it is a meeting to study for a class, the system should allow the students to be able to quickly pull up notes from the class and automatically associate it to that class for each student at the meeting.

Perhaps by now the reader has also raised some cautionary flags on the societal impact of ubiquitous computing, at least in the form that we have presented here. Certainly there are issues to do with privacy and security that cannot simply be ignored in the drive to push technology off the desktop and into the very fabric of our lives. We must and will consider how the benefits of our ubiquitous computing applications weigh up against the threat of Big Brother [47]. Again, we do not know currently how to design systems that reduce, not increase, this threat. It is another serious research effort that must be considered as a design and evaluation challenge that requires an inherently interdisciplinary approach. Such interdisciplinary research is a foundation for many researchers in Georgia Tech's GVU Center. Perhaps more qualitative and naturalistic evaluation methods might uncover social miscues with ubiquitous computing better than more empirical and positivist techniques.

## 2.4 Prior Work

The concept of ubiquitous computing was initially defined by Mark Weiser in 1991 [67]. According to Weiser, both the ubiquity and the transparency of the system are the salient characteristics of ubiquitous computing. There are two main approaches to satisfy the ubiquity requirement: either the system is mobile enough to be carried by the user, or the system is an instrumented space that surrounds the user in her/his location, e.g., a room, an office, a building. (provides a surrounding perspective of the users). These two approaches actually complement each other and many ubiquitous computing installations rely on a synergy of instrumented spaces and mobile devices.

Here, we briefly outline the state of the art in the area of ubiquitous computing as it it applies to a) providing access to computational services, b) making the computational devices transparent to the users by embedding them into the environment, and c) generating, intelligent, interactive and reactive systems. All of this work on ubiquitous computing has a strong motivational impact on our interests in establishing a large-scale ubiquitous computing Infrastructure.

Many ubiquitous systems rely on a fixed infrastructure to provide computational services to the user. Perhaps one of the earlier incarnations of such a system is the Olivetti Active Badge system [64]. It relies on an infrastructure of infrared beacons installed in every office and room throughout a building. Users wear a pager-sized transponder device that identifies them and communicates their location to a central server. Some of the initial applications of this technology were to quickly locate a person on a campus or "teleporting" automatically their usual desktop environment to the closest workstation [55].

The Xerox ParcTab system [65] goes one step further in augmenting the user with mobile computational power. ParcTabs, in addition to having active badge style location awareness are also fitted with a monochrome screen and a pen-based interface. A communication infrastructure is also added to provide mobile users with access to their workstations or servers data: remote file browsing, dictionary, weather forecast, web access, etc.

A more recent extension to Active Badge systems is the Audio Aura system, in which the users wear an active badge and wireless headphones [46]. When they move in the building, RF transmitters deliver personalized information according to the user's location in form of an audio signal. Several other extensions included the Palplates system, which consists of small fixed touch-sensitive displays installed in strategic locations in an office environment: in meeting rooms, in the kitchen, next to the printers [39].

In all the above described systems, the emphasis is on ubiquity: computational services are available wherever the user is. But access to these services is hardly transparent: users have to wear or operate devices. Audio Aura is the most transparent of the systems above, except for the fact that users have to wear headphones which isolate them from the surrounding world. Significant contributions in the area of designing artifacts that make the interface between the user and the computational resources transparent are presented by Ishii [30]. Similar tendencies, however, much more concentrated on embedding the computing, intelligence and responsiveness in devices is the charter of the Things that Think Project at the MIT Media Laboratory [44].

Augmented reality has much in common with ubiquitous computing. An early augmented reality system that let the user access computational power through real world objects is Wellner's Digital Desk [69]. This system replaces the traditional screen-based desktop metaphor with an actual desktop. A ceiling-mounted camera and projector allow the desktop surface to be used as an input/output device. The user can draw on an actual paper sheet and benefit from computational services like copy/paste (an actual line drawing may be copied and is then pasted: the copies are actually projected.) The user can also start from an existing digital document (i.e., projected) and use real-world tools like pen, eraser or even her hands to manipulate data. The real-world desktop is *augmented* with computational capabilities. Mackay's Ariel prototype moves this interaction technique into the field for mobile engineers annotating engineering drawings on a construction site [37].

There are also other augmented reality research prototypes that require the user to wear some apparatus like see-through displays that project annotations onto a real-world view. However, users are free to move around in space (e.g., a library, a maintenance room). Feiner's Karma system for the maintenance of copy machines [24] and Boeing's wearable system for aircraft maintenance [60] are representative of this approach. Extensions of this type of work have resulted in a newer field of research entitled Wearable Computing [66]. The interests of research in this area is towards human augmentation by embedding the computing into clothing or artifacts of daily wear.

In a similar vein, research on home automation has focused on hiding computational devices for aesthetic reasons and providing transparent interaction to accommodate non-technical users. Research efforts at Microsoft Research or IBM T.J.Watson Research Center are representative of this trend. The most interactionally-transparent home research effort to-date is the Neural Network House [45]. Fitted with sensors to detect people's presence, motion and actions, the system regulates lights and heating in a house according to patterns of usage. Transparent interaction has also been explored for office environments [63]. A well developed example of transparent interaction in a conference room is The Reactive Room [13].

More recently, there has been a lot of interest in even pushing deeper towards having interfaces that are completely unencumbered and non-invasive. Computer vision and audition techniques are being extensively employed to provide ubiquitous systems with awareness of the users. Vision systems are being used to locate users in the environments, recognize them, and track their gestures, expressions, and body poses. The Reactive Room [13] also uses computer vision techniques to track users and so do the Smart Rooms [50] and KidsRoom [8] systems from the MIT Media Laboratory. Other instances of smart and interactive environments that use computer vision techniques to track and recognize human activity are being developed at the MIT AI Laboratory [43, 42], CMU [29, 57], University of Maryland [27, 71] and several other leading research institutions.

## 2.5 Research Training Activities and Education Involving The System

We expect the development of this equipment will involve a substantial number of graduate and undergraduate students, and be the subject of many PhD and undergraduate theses. Research training is an integral part of our research activities, in which students learn by doing as well as explicit instruction. We have a long and successful history of tightly involving graduate and undergraduate students in our research projects, and many have gone on to work at places like Xerox PARC, IBM, Netscape, Fuji-Xerox, MIT Media Lab, and Berkeley.

In addition, a major focus of our research is how ubiquitous interfaces can be used in education (the Classroom 2000 project, for example). We hope to revolutionize the classroom experience, which of course would have a tremendous impact on research training.

### 2.6 Who Would Use The System?

The personnel developing the instrumentation include the 2 budgeted postdoctoral fellows, the 4 PI faculty, several additional faculty, and approximately 5 graduate students and 10 undergraduates. The personnel using the instrumentation to explore issues in ubiquitous and aware computing include the PIs and additional faculty faculty, an additional 2 postdoctoral fellows, 2 or more visiting faculty, and approximately 10 graduate students and 20 undergraduates. We expect a large fraction of the College of Computing to use this system in their daily lives, including almost every class taught in the College of Computing. The College consists of 46 faculty, 17 postdocs and research staff, 180 graduate students, and 1200 undergraduates.

## 3 Description of the Research Instrumentation and Needs

### 3.1 Existing Equipment and Infrastructure

We have access to computational and network capabilities from a variety of units across the Georgia Tech campus, including the College of Computing, the GVU Center, the Broadband Telecommunications Center (BTC), the Office of Information Technology (OIT), and the Scientific Visualization Lab. We will first describe the general computational and networking infrastructure of these units and then focus on specialized infrastructure that directly supports our efforts in ubiquitous computing applications.

The College of Computing has more than 350 workstation-class machines from Sun, Silicon Graphics, Intel, Apple, Hewlett-Packard, and IBM; and over 150 X-terminals. To support these workstations, we have 40 Sun and SGI systems used as file and compute servers, 10 of which are quad-processor machines. Cluster parallel systems include a cluster of 16 Sun UltraSPARC processors and 5 UltraSPARC dual-processors utilizing Myrinet (8 systems), Dolphin (4 systems), and ATM interconnects (all systems), a cluster of 16 Pentium quad-processors interconnected by Myrinet and 100 Megabit Ethernet, an SGI Origin 2000 quad-processor, an SGI Power Challenge with 12 processors, (these two are expected to be upgraded to a 16 processor SGI Origin 2000), a cluster of 16 Silicon Graphics R4400 processors utilizing an ATM interconnect, SGI, Sun, and Intel video servers, and high-end SGI, Sun, and Intel multimedia workstations. In addition, the campus-wide high-performance machines including an IBM SP-2 and a 20-processor SGI Origin 2000 Multiprocessor are available for experimentation. All of these facilities are linked by a dedicated high-performance network utilizing ATM.

All of the College's facilities are linked via local area networks that can provide a choice of switched 10 Mbps Ethernet, switched or shared 100 Mbps Ethernet, 155 Mbps ATM, or 622 Mbps ATM connections to most locations, including offices, labs, and classrooms. The College's network employs a high-performance OC12C ATM (622 Mbps) and FDDI (100 Mbps) backbone with connectivity to the campus ATM network and to the global Internet beyond via redundant OC3C ATM (155 Mbps) links. The primary campus Internet connection is provided by a direct FDDI (100 Mbps) link to the service provider's Atlanta switching center. An OC3C (155 Mbps) ATM connection to the NSF vBNS (very high performance Backbone Network Service) research network has recently been brought into service. Georgia Tech is also leading efforts to establish a southern regional gigabit network as part of Internet2.

The Graphics, Visualization, and Usability (GVU) Center houses a variety of graphics and multimedia equipment, including high-performance systems from Silicon Graphics, Sun, Hewlett-Packard, Digital, Intel, and Apple, as well as extensive video and audio facilities for recording and editing. The Scientific Visualization Laboratory has additional equipment from Digital and Silicon Graphics. The Broadband Telecommunications Center (BTC), have several labs equipped with leading edge computing, communications, and test equipment. These include the Hybrid Fiber/Coax (HFC), Asynchronous Transfer Mode (ATM), Home Information Infrastructure (HII), Protocols, Wireless Technologies, and Video Sources Labs.

The GVU Center and College of Computing opened an auxiliary lab in September 1997 that contains a 3000 square foot facility with 50 workstations (Macintosh, PC, SGI, Solaris) and specialized equipment for robotics, ubiquitous computing and virtual environments. This specialized equipment consists of a number of LCD projectors and electronic whiteboards, that will be directly used in this project. In addition, this unit has the beginnings of a commercial, high-speed wireless Ethernet network for use with mobile computing platforms.

In January 1997, the College of Computing and the Office of Information Technology opened the prototype classroom that is the principal site for Classroom 2000 evaluation, described below. It is our experience in designing, implementing and using this prototype classroom, described below, that is the main impetus for this more ambitious instrumentation proposal. As a result of the success of Classroom 2000, the College received an additional $50,000 from the university to replicate essential features of the prototype classroom in another classroom.

## 3.2 Preliminary Results From Existing Equipment

In the past 2 years, the Future Computing Environments (FCE) group at the College of Computing at Georgia Tech has developed a number of applications that rely on the concepts of ubiquitous and aware computing. This work includes the project *Classroom 2000,* in which we instrumented a single room [3, 10, 4]. The *Domisilica* project explores ubiquitous and aware computing in the home and is a joint project with the BTC on campus [38, 40, 41]. Our *Cyberguide* project provided handheld interfaces for exploring indoor and outdoor spaces on the campus of Georgia Tech [35, 36, 5]. We describe these different examples below. More information can be found at http://www.cc.gatech.edu/fce.

**The Classroom:** As outlined in Section 2.3.1, we built a special-purpose instrumented classroom to support live experimentation with Classroom 2000. This classroom has specialized presentation and capture facilities suited for the Classroom 2000 project. Specifically, there is a 67-inch diagonal rear-screen projected Liveboard and two additional ceiling-mounted LCD projectors from Proxima Corporation. There are high-quality ceiling microphones and wireless lapel microphones for recording of audio signals in the room. Ceiling-mounted cameras provide several video feeds. All audio/video feeds are coordinated by a mixer that records on both analog and digital media. The periphery of the room contains 15 teacher/student workstations (Macintosh/PC/SGI/Solaris). Each of 40 students seats is provided with dedicated power and Ethernet through a raised floor. We have not yet had the funds to put adequate machines in the hands of the students at each of these locations, but with expected donations from Hewlett-Packard (see letter that accompanies this proposal), we will soon have this.

Over the past two years, Classroom 2000 has been used regularly in 25 different undergraduate and graduate courses. The system is popular with students and professors and its use is increasing. We have built up a substantial database of captured material (see www.cc.gatech.edu/fce/c2000). The prototype classroom is being replicated across the Georgia Tech campus, and at other universities in Georgia (Kennesaw State University and Georgia State) and elsewhere (University of California, Berkeley and the University of Michigan). Section 4 discusses other universities and companies using our Classroom 2000 software. Our preliminary results have helped us to form a long-term vision of how ubiquitous computing can improve the educational experience [2, 4, 3, 10]. Our goals are in alignment with those expressed in a recent NSF workshop report on the role of Computer Science in educational technology [28].

The current implementation of Classroom 2000 has been instrumented to give us a better idea of how the system performs and is used in everyday classroom activity. These measurements are at the systems level as well as at the application level. The majority of our measurements have been aimed at developing a better picture of end-user access patterns, and have been initially reported in [4, 3]. The feedback we are getting from actual users of the system has proved invaluable and has lead directly to the inclusion of capabilities such as an intelligent login, automatic synchronization for video/audio recording, capture of Web activity, and more informative and available interfaces for accessing augmented class notes.

The systems group at Georgia Tech has a long history of research in parallel and distributed systems relevant to building the computational infrastructure for large-scale distributed systems such as the ubiquitous interface. Research results that are of particular relevance to this research include (a) the design and development of object-based distributed operating systems [14], (b) the development of consistency models for maintaining shared state in parallel and distributed systems [52, 34, 61, 6], and (c) efficient implementations of distributed shared memory systems on cluster machines [53, 61, 62, 31]. In an attempt to provide benchmark statistics to feed a scalability simulation for future media server capabilities (see 2.3.4), we have also performed measurements on the traffic between the electronic whiteboard and a central server machine (a uniprocessor SunSparc workstation). In a typical course, the average pen stroke data rate is 2-300 KB/sec. Our immediate plan is to further instrument the current system to measure the data rates (average, max) between capture clients and the server, and between the server and the file system.

**The Home:** The Domisilica project is aimed at producing a virtual community that mirrors and supports a real physical community. Our initial efforts are targeted toward the home and the extended family. We have built a prototype virtual home environment that is tied to a home setting of a number of researchers in FCE. We are making the two worlds, physical and virtual, work in concert with each other. So, for example, when some produce is placed inside the physical refrigerator in the kitchen, the contents of a virtual refrigerator, CyberFridge, is automatically updated as well. We are also experimenting with how activity in the virtual world can affect the physical world. For example, when multiple people virtually visit a room in Domisilica that is associated with a physical room, say a living room, the physical environment produces more ambient noise to inform the physical occupants of the room of the presence of the virtual visitors. We are interested in developing more automatic ways to communicate between the virtual and the real worlds. This will be achieved by adding sensors to the environment that will identify the user and the activity in the environment and update the virtual representation of this environment.

**Personal space:** The previous two examples dealt with fairly well defined physical spaces, the classroom and the home. We are also interested in pursuing the concepts of ubiquitous and aware computing in environments where the physical space is defined as the unfamiliar territory that surrounds a mobile user. The Cyberguide project is aimed at developing mobile assistants, more specifically, tour guides that are aware of the location and orientation of their user and provide information about the surrounding space. Our initial work in this area has relied on using different forms of hand-held computers with position sensors. So far we are concentrating more on the software development issues with the hope of keeping it platform independent. We are also pursuing some research on wearable computers within this context.

### 3.2.1 Why we need new equipment and infrastructure

With all of this existing infrastructure, it is reasonable to ask whether additional instrumentation is required for the research activities outlined in this proposal. Some equipment is not available at Georgia Tech, and for other equipment exclusive access is required. Our experimental facility requires sensors and display equipment to be installed that we do not have at this time. We need new wiring and wireless networking to support the connection between interface devices and central servers. Although we have a substantial amount of computational power available, this power is also used for other purposes. We need dedicated computational and storage resources to support full and continuous availability of the ubiquitous interface system.

## 3.3 System Design

The concept behind our intended instrumentation design is to identify different types of spaces and then instrument those spaces according to their needs. The spaces we will consider for instrumentation are:

- **Large lecture-style classrooms.** We envision supporting two of these, all in the College of Computing, in addition to the original Classroom 2000 prototype.
- **Small formal meeting/seminar rooms.** There will be six of these spaces distributed across all buildings.
- **Social spaces with tables and chairs, or an informal meeting space.** There will be one social space on each floor of each building. We will develop a modular set-up for these spaces and will deploy 9 of them across the buildings on campus.
- **Walk-around public spaces, such as foyers to buildings and hallways.** All hallways surrounding classrooms, meeting rooms and social spaces will be considered as well as one entry way for each floor of a building. The centerpiece of these spaces will be large, interactive bulletin boards, and we will provide 5 of these at critical entryways to the various building floors.
- **Private and semi-private offices.** There will be a small number of these in different buildings. They will be configured similarly to the social spaces with tables, but will have very different uses, of course. We will build three of these offices.

The instrumentation of these systems falls into several major categories:

- **Interactive displays.** This category consists of large-screen electronic whiteboards, gas plasma flat screen video walls as well as portable laptops with tablet inputs.
- **Audio and video sensing.** One of the essential features of the ubiquitous computing environment is the ability to sense audio and video signals. This is for detecting who is present and what is happening, and for capturing what occurs in the surrounding space for later access, as is done in Classroom 2000.
- **Computational infrastructure.** Every video camera will need local computation to support perception tasks.
- **Network infrastructure.** Audio will be captured and delivered via a commercial Ethernet transport system (the RAVE system from DSQAudio Systems). In addition, we will need to increase the capacity of the fiber ATM campus network that connects three of the buildings under consideration.

We will now provide further details on instrumentation for the separate physical spaces. During the actual grant, we will stage the deployment of the various spaces, spending the first year designing, developing, and installing one prototype for each space. The subsequent two years will be spent scaling the ubiquitous space up, to include all proposed spaces.

### 3.3.1 Instrumenting a classroom

The classroom will contain a large electronic whiteboard consisting of three rear-projected SmartBoard RP 1602 units from Smart Technology, Inc. Each unit is a 58-inch diagonal screen with a separate networked computer and LCD projector. We have developed a software system, ZenPad, which will coordinate the displays between these three boards to make them act as a coordinated unit for the lecturer. The particular advantage of the SmartBoard is that the networked computer can be a Macintosh, PC or UNIX box, allowing for the greatest flexibility for instructors wanting to do on-line demonstrations. We have priced the electronic whiteboard at $7,524, with educational discount. The projectors will be Proxima 9200s, priced at $8500. We will begin with Pentium II 300MHz PCs for each board, each priced at $3,500 each.

Students will be seated two to a table. Furniture for Classroom 2000 cost $300 per table and $150 per chair. We will configure the classroom to hold 40 students. Each student will have access to a laptop computer (these are being donated by Hewlett-Packard). We will attach a Wacom ArtZII 6x8 digitizing tablet to provide an economical pen-based input for taking electronic notes. Again, the software infrastructure for ZenPad was designed to support student note-taking as well as that of the lecturer. The in-class networking infrastructure will be provided by a 5.8GHz wireless network from RadioLAN, rated at 10 Mbps. Several access nodes/bridges for the network will reside in the classroom, at a projected cost of $8085 and each interface card for the student laptops will cost $500.

Audio will be recorded with an array of networked microphones. The lecturer will wear a wireless lapel microphone (we have used an AudioTechnica 6100, priced at $500). Each student table will have a single high-quality microphone ($100 each). Nine video cameras will be mounted in the ceiling of the room in a regular array pattern. The eight perimeter cameras will be Sony EVI/D30 pan/tilt/zoom cameras ($1000 each) and the middle camera will be an OmniCam, a soon-to-be-commercialized panoramic camera from the University of Columbia (estimated price is $1200). Each camera will be attached to a Dual Pentium II computer ($5500) that will capture the video signal as well as remotely control the camera via serial cable.

There is other specialized equipment for the classroom. Scanners are being donated by Hewlett-Packard. Servers to host the captured classroom material are being donated by Sun Microsystems. The networked audio will be collected in a central machine room/audio center where it will be fed to a digital mixer (a Tascam 48-channel digital mixer costing $10,000). The audio for the room will be on a dedicated Fast Ethernet network, at a cost of $200 per node/microphone.

In total, the classrooms will be augmented at an estimated cost of $214,612 per room.

### 3.3.2 Instrumenting a formal/private meeting room

In this space, we assume there is a single conference table that will accommodate 10 people. We will provide a ceiling-mounted Proxima 9200 projector ($8500 with $1000 for ceiling installation) and a SmartBoard 380 wall-mounted digitizing whiteboard ($2,099 with educational discount) with a 72-inch diagonal interactive surface. A Pentium II class machine ($3500) and SVHS projector ($400) will be attached to the projector. A Polycom star teleconferencing microphone $999) will provide 4 separate microphone inputs. A sound reinforcing system will be built into the room at an estimated cost of $6000. We will also provide 4 note-taking devices similar to those described in the classroom. There will be 4 ceiling mounted Sony EVI/D30 pan/tilt/zoom cameras and a single Omnicam mounted on the table or ceiling. Each camera will be connected to a Dual Pentium class machine for local processing and interim storage.

In total, each formal meeting room will be augmented at an estimated cost of $59,498 per room.

### 3.3.3 Instrumenting an informal meeting space

This space consists of a number of tables/booths where students meet to eat lunch, socialize, study and otherwise work. We are proposing a modular system to support each table/booth. Each module will provide a table-mounted Polycom star teleconferencing unit ($999) with $500 spent on a sound reinforcement and audio isolation system. A single Sony EVI/D30 pan/tilt/zoom camera ($1000) will be mounted overhead and an Omnicam ($1200) will be table-mounted. Each camera will be attached to a Dual Pentium II class machine ($5500 each). A single note-taking laptop will be networked and tethered to the table.

In total, each module of this informal meeting space will be built at a cost of $15,499.

#### 3.3.4 Instrumenting private offices

Though the principal focus of this instrumentation is not for private spaces, we will nonetheless strive to instrument a small number of private offices to experiment with personalized access to a ubiquitous campus interface. The offices will be provided with exactly the same module infrastructure as described for informal meeting places, at a cost of $15,499 each.

#### 3.3.5 Instrumenting a walk-around public space

In this space, there are no places to sit down. This is a hallway or an entrance area to a building or section of a building. People walk around this space and we want to be able to detect their presence. We will use Toshiba 1k-627AT 1/3-inch CCD security cameras (20 at $500 each). We also want to provide large-scale interactive kiosk-like displays that will serve as interactive bulletin boards. These bulletin boards will consist of a Philips 42-inch Flat-TV plasma display with 2 Elmo 1/4-inch QN42H CCD color cameras ($2385 each with controller card) and a Dual Pentium II class machine ($5500). We will deploy 5 of these interactive bulletin boards throughout our buildings.

#### 3.3.6 Other required infrastructure

We will use a series of DSQAudio RAVE Audio Ethernet Routers to provide a dedicated Fast Ethernet network for up to 64 microphones. Each classroom will have its own dedicated network that will consist of 3 16-port Model 161 routers ($4000 each). The social spaces, private offices and small meeting rooms will be served by 9 Model 188 routers ($3900 each). A single machine room will be located in two buildings providing digital mixing (Tascam 48-channel mixers for the classroom at $10,000 each and Yamaha O3D 16-channel digital mixer for the remaining microphones at $3699 each). To support inter-building networking, we will need to augment the existing ATM OC12 fiber backbone with an estimated $200,000 of additional equipment (the price breakdown of this networking infrastructure can be provided on demand).

The core networking infrastructure supporting the project consists of three identical Cisco Catalyst 5500 ATM/LAN switches, one for each of the principal sites (CCB, GCATT, CRB), connected in a fully meshed topology via OC-12C (622 mbps) links over existing dark fiber. Within each of these buildings, the switches provide 4 OC-3C (155 mbps) ATM ports for connection of related servers and 24 switched 100BaseT (100 mbps) Fast Ethernet ports, with cell-to-frame bridging between the ATM and Ethernet switching fabrics. Additional ATM and LAN slots are provided for future expansion. At the time of actual purchase, if there is alternate technology supplying greater bandwidth with similar quality of service (QOS) capabilities, it will of course be considered.

#### 3.3.7 Required personnel

We will be hiring two post-docs to help lead some of the research activities involved with the design, implementation, and evaluation of the various ubiquitous computing applications made possible by this instrumentation grant. At least one of these researchers will be responsible for some of the daily operational matters of this large-scale system. Matching money from the university will allow us to also hire a part-time technical support staff person.

We have learned from experience that a critical factor in deploying such a ubiquitous computing environment requires expertise that is most economically acquired in the form of outside consultants. Therefore, we are budgeting to use the remaining university matching money to pay for design consultants to assist us in finalizing plans for various spaces. We estimate that such professional design consulting services will cost us between $25,000–40,000 over the course of the three year grant.

#### 3.3.8 Use of off the shelf software

We are using and plan to expand our use of several commercial speech recognition products, including the HTK toolkit from Entropic [18] and Naturally Speaking from Dragon Systems [17]. Neither system is perfect, but we have been able to get usable results in pilot studies. One experimental question we would like to address is how good speech recognition has to be to support functions like searching for relevant sections of an audio stream. We suspect that for many tasks current speech recognition performance is usable. In the Classroom 2000 system the output of the speech recognition module is a sequence of time-stamped tokens. In cases where there is possible confusion multiple interpretations are put in the token stream with confidence levels associated with each interpretation. Later processes use the multiple interpretations in different ways. A module that generates a transcript for a human reader picks the most likely interpretation. A module which supports keyword search can search all interpretations for a match.

On the system software side, we plan to leverage heavily off both commercially available packages, as well as experimental software subsystems developed both locally at Georgia Tech and elsewhere to realize the runtime infrastructure. Specifically, we will use experimental[48] or custom[56] communication software to enable low latency, high throughput, and time sensitive communications within the clusters. Communication between the clients and the media server will be layered on top of IP to ensure ubiquitous access to the media server. We plan to use the system technologies (such as space-time memory and cluster-wide threads support, developed jointly with DEC CRL) from the Stampede project as a starting point in our prototype development. We will use standard file system interfaces to implement the data repository.

#### 3.3.9 Why use distributed servers?

A cluster of commodity boxes interconnected by high speed networks is an ideal vehicle for achieving the scalable parallelism needed in our system. We expect many users to be simultaneously accessing the system. With the increase in the computational capacity of desktop computers, and the emergence of low-latency high-bandwidth cluster interconnects, a network of commodity boxes (workstations and/or PCs) is emerging as a cost-effective parallel processing platform. The boxes may themselves be small-scale SMPs or uniprocessors. We have a significant amount of experience in building cluster servers for specific applications both at Georgia Tech and in association with companies such as DEC [62, 54]. We expect that a large-scale media server for applications such as this system to be implemented on multiple interconnected clusters.

## 4 Impact of Infrastructure Project

The purpose of this proposal is to create a unique testbed for the development and exploration of large-scale applications of ubiquitous computing. The existence of this experimental facility will play a major role in developing a new approach to computing. We expect that ubiquitous and aware computing will play a dominant role in defining how we think about computing in the near future. The new approach represents a shift from thinking about computers as extensions of the desktop to extensions of the environment we live in. Experimental results and associated demonstrations will play a critical role in this paradigm shift. Small-scale implementations, typically limited to a single chair or room, are useful but do not support exploration of large-scale applications.

In terms of research this grant will allow the College of Computing and the Graphics, Visualization, and Usability Center at Georgia Tech to become a center of excellence in research on ubiquitous and aware interfaces. In terms of education a major focus of our research is how ubiquitous interfaces can be used in education (the Classroom 2000 project). It is also important that the many students who study computer interfaces in the College of Computing and the Graphics, Visualization, and Usability Center are exposed to new interface paradigms. We have recently inaugurated a Master's program in Human Computer Interaction.

We also have a long and successful history of tightly involving graduate and undergraduate students in our research projects, and many have gone on to work at places like Xerox PARC, IBM, Netscape, Fuji-Xerox, MIT Media Lab, and Berkeley. Five graduate students (2 women) and 20 undergraduates (1 woman, 2 minority) have worked extensively on Classroom 2000, and an additional 40 graduate students (10 women, 3 minority) have used the system for projects of various sorts. We expect the ubiquitous interface project to be a major magnet for students of all levels and of all kinds: we expect more undergraduates will get excited about computing and want to go to graduate school, and more graduate students will do exciting and innovative theses.

Our work in ubiquitous computing has already had impact at other universities. We work hard to make sure useful data from our work is available on the World Wide Web (www.cc.gatech.edu/fce/c2000). Versions of the Classroom 2000 system have been installed at other universities outside of Georgia Tech, including Kennesaw State University, University of California, Berkeley and the University of Michigan. Plans are underway to install Classroom 2000 at Georgia State University. Several industrial research labs, including Smart Technologies, Microsoft Research, BellSouth and the Corporation for National Research Initiatives, are planning to evaluate extensions of the Classroom 2000 infrastructure to support non-educational applications. We expect this pattern of technology transfer to continue. We also expect other institutions to want to copy or improve on our experimental facility.

We strongly feel this project will "expand the scope of research and research training", "foster the integration of research and education by providing instrumentation for research-intensive learning environments", and "foster the development of the next generation of instrumentation for research and research training." The research supported by this experimental facility advances the goals articulated in "NSF In A Changing World," the Foundation's strategic plan (NSF 95-24). We expect to develop new approaches to education and workforce training and machine assistance and augmentation that can potentially reach all Americans. We expect to directly explore potential roles of information infrastructure in many aspects of daily life. Our goal is the development of an intelligent, agile, and adaptable infrastructure for the future – one that takes full advantage of the capabilities of the emerging information infrastructure.

The intellectual merit of the proposed activity is shown in a number of ways. The system we will build is critical to advancing knowledge and understanding in the exploration of new ways to interact with computers. The PIs who will develop this instrumentation have already developed small versions of the system, and shown they can build working systems that others can learn from, use, and replicate at their own institutions. Although the ideas of ubiquitous and aware interfaces have a long history, especially in the science fiction literature, the creativity and originality of the concepts involved in this research are demonstrated by the fact that this facility will be unique. We hope our description of the system convinces you that it is well conceived and organized. The system is appropriate and required for our current and expected research and training in human–computer interfaces. There are sufficient institutional resources to maintain and operate the system.

The development of this system will advance discovery and understanding while promoting teaching, training, and learning. In fact, a major use of the system is to explore possible roles of ubiquitous computing in education. We intend to participate in an existing College of Computing summer internship program for minority undergraduates from other institutions, and we will encourage these students to work with us. The GVU center has an NSF Traineeship, which we use to aggressively recruit women and minorities. The proposed system will be designed so as to complement existing computational infrastructure in the College of Computing and at Georgia Tech. We will work hard to disseminate our results broadly, both by publishing papers, by demonstrating the system, and by encouraging visiting researchers. The benefits of this research to society include having an experimental facility which can explore and understand the consequences of possible future technology, before that technology is already widespread.

The construction of this experimental facility will promote a partnership between the College of Computing and private sector instrument developers such as Sun, Hewlett-Packard, BellSouth, and Motorola. We will work with experts from these companies to develop an effective system, and we and they expect there will be many new product spinoffs.

# 5 Project and Management Plan

The goals of our project and management plan are to: 1) Ensure proper administration of the funds for the purchase, support, and maintenance of the requested instrumentation. 2) Formulate policies to decide who can use and/or modify the system and when. 3) Facilitate communication with other research groups in academia and industry. 4) Influence industry and make technology transfer easier. 5) Make others in academia aware of our work and facilitate technology and idea transfer among universities working in this area. The four PIs (Abowd, Atkeson, Essa, and Ramachandran) and the two postdocs will form the management committee for the system. This group will make day to day decisions about policies and use of funds. We will also have an advisory group made up of researchers from universities and industry. We plan to invite the following individuals to be initial members of the advisory group: Rob Bouzon (Hewlett-Packard), John Hale (Sun Microsystems), Ron Borgstahl (Motorola), Lee Friedman (BellSouth), Ira Richer (Corporation for National Research Initiatives), Jim Foley (Mitsubishi Electric Research Labs), Bill Schilit (Fuji Xerox Palo Alto Labs), Elizabeth Mynatt (Xerox PARC), Mark Weiser (Xerox PARC), James Landay (University of California, Berkeley), Peter Freeman (Dean, Georgia Tech College of Computing), John Limb (Director, Georgia Tech Broadband Telecommunications Center), and Jarek Rossignac (Director, Georgia Tech GVU Center). The function of the advisory group will be to provide strategic guidance and to facilitate communication and dissemination of our work.

*Allocating instrument time:* Our goal is to encourage use of this system as much as possible. The experimental system will be spread throughout the College of Computing. It will be possible to subdivide the system to explore different algorithms in different places. It will also be possible to run different algorithms at different times. If it becomes necessary to arbitrate competing usage requests the management committee will decide how to handle them. No user fees are planned.

*Maintenance, operation, and technical support:* We have included two postdoctoral fellows in the budget, who will coordinate maintenance, operation, and technical support. Maintenance costs will be provided from the materials and services portion of the budget, as well as from other sources. It is difficult to estimate repair costs for fabricated equipment, but we hope to minimize these costs.

*Attracting new users.* We have a track record of successfully transferring complex software to other institutions. We expect to continue that type of technology transfer for this system. In addition, we expect to set up a visitor program that will facilitate short and long term visits by other investigators and students. We will apply for NSF REU support to sponsor many undergraduates for summer internships. We have an existing summer internship program for minority undergraduates from other institutions, and we will encourage these students to work with us.

In addition a weekly research seminar on ubiquitous and aware computing will encourage communication among the researchers and users of the system.

## References


[1] Gregory D. Abowd. Ubiquitous computing: Research themes and open issues from an applications perspective. Technical Report GIT-GVU 96-24, GVU Center, Georgia Institute of Technology, October 1996.

[2] Gregory D. Abowd, Chris Atkeson, Ami Feinstein, Yusuf Goolamabbas, Cindy Hmelo, Scott Register, Nitin "Nick" Sawhney, and Mikiya Tani. Classroom 2000: Enhancing classroom interaction and review. Technical Report GIT–96–21, Graphics, Visualization and Usability Center, Georgia Institute of Technology, September 1996.

[3] Gregory D. Abowd, Christopher G. Atkeson, Jason Brotherton, Tommy Enqvist, Paul Gulley, and Johan LeMon. Investigating the capture, integration and access problem of ubiquitous computing in an educational setting. In *Proceedings of the 1998 conference on Human Factors in Computing Systems — CHI'98*, 1998. To appear.

[4] Gregory D. Abowd, Christopher G. Atkeson, Ami Feinstein, Cindy Hmelo, Rob Kooper, Sue Long, Nick Sawhney, and Mikiya Tani. Teaching and learning as multimedia authoring: The Classroom 2000 project. In *Proceedings of the ACM Conference on Multimedia — Multimedia'96*, pages 187–198, 1996.

[5] Gregory D. Abowd, Christopher G. Atkeson, Jason Hong, Sue Long, Rob Kooper, and Mike Pinkerton. Cyberguide: A mobile context-aware tour guide. *ACM Wireless Networks*, 3, 1997. To appear.

[6] Mustaque Ahamad, Gil Neiger, James E. Burns, Phillip W. Hutto, and Prince Kohli. Causal memory: Definitions, implementations and programming. *Distributed Computing*, 1995.

[7] Suzanna Becker and Mark Plumbley. Unsupervised neural network learning procedures for feature extraction and classification. *Journal of Applied Intelligence*, 6:1–21, 1996.

[8] A. Bobick, S. Intille, J. Davis, F. Baird, C. Pinhanez, L. Campbell, Y. Ivanov, A Schutte, and A. Wilson. The KidsRoom: a perceptually-based interactive and immersive story environment. Technical Report 398, MIT Media Laboratory, Perceptual Computing Section, 1997. Available from `http://www-white.media.mit.edu/vismod/`.

[9] M. Brand and I. Essa. Causal analysis for visual gesture understanding. Technical Report 327, Perceptual Computing Group, Media Laboratory, M.I.T., 1995. Appeared in AAAI Fall Symposium on Computational Models for Integrating Language and Vision. Available as MIT Media Lab Perceptual Computing Techreport # 327 from `http://www-white.media.mit.edu/vismod/`.

[10] Jason A. Brotherton and Gregory D. Abowd. Rooms take note! Room takes notes! In *Proceedings of the 1998 Spring AAAI Symposium on Intelligent Environments*, 1998.

[11] Claire Cardie. Emperical methods in information extraction. *AI Magazine*, 1998. to appear.

[12] Chris Clack, Jonny Farrington, Peter Lidwell, and Tina Yu. Autonomous document classification for business. In *Proceedings of the 1997 International Conference on Autonomous Agents*, February 1997.

[13] Jeremy R. Cooperstock, Koichiro Tanikoshi, Garry Beirne, Tracy Narine, and William Buxton. Evolution of a reactive environment. In *Proceedings of the ACM SIGCHI'95 Conference*, 1995.

[14] Partha Dasgupta, Richard LeBlanc, Mustaque Ahamad, and Umakishore Ramachandran. The clouds distributed operating system. *IEEE Computer*, 24(11):34–44, November 1991.

[15] Scott Deerwester, Susan T. Dumais, and Richard Harshman. Indexing by latent semantic indexing. *Journal of the American Society for Information Science*, 41(6):391–407, 1990.

[16] Anind Dey, Gregory D. Abowd, and Andy Wood. Cyberdesk: A framework for providing self-integrating context-aware services. In *Proceedings of the 1998 Intelligent User Interfaces Conference — IUI'98*, pages 47–54, 1998.

[17] DragonDictate and Dragon Naturally Speaking, Speech Recognition products by Dragon Systems Inc. `www.dragonsystems.com`.

[18] HTK, ESPS, WAVES+, GrapHvite and other Speech and HMM related products by Entropics Research Laboratories Inc. `www.entropic.com`.

[19] I. Essa, S. Basu, T. Darrell, and A. Pentland. Modeling, tracking and interactive animation of faces and heads using input from video. In *Proceedings of Computer Animation Conference 1996*, pages 68–79. IEEE Computer Society, June 1996.

[20] I. Essa, T. Darrell, and A. Pentland. Tracking facial motion. In *Proceedings of the Workshop on Motion of Nonrigid and Articulated Objects*, pages 36–42. IEEE Computer Society, 1994.

[21] I. Essa and A. Pentland. A vision system for observing and extracting facial action parameters. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 76–83. IEEE Computer Society, 1994.

[22] I. Essa and A. Pentland. Facial expression recognition using a dynamic model and motion energy. In *Proceedings of the International Conference on Computer Vision*, pages 360–367. IEEE Computer Society, Cambridge, MA, 1995.

[23] I. Essa and A. Pentland. Coding, analysis, interpretation, and recognition of facial expressions. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 19(7):757–763, 1997.

[24] S. Feiner, B. MacIntyre, and D. Seligmann. Knowledge-based augmented reality. *Communications of the ACM*, 36(7):53–62, 1993.

[25] David J. Finton and Yu Hen Hu. Importance-based feature extraction for reinforcement learning. In et al. Thomas Petsche, editor, *Computational Learning Theory and Natural Learning Systems*, volume 3. MIT Press, 1994.

[26] National Science Foundation. NSF In A Changing World (NSF strategic plan, NSF 95-24), 1995. from http://www.nsf.gov/cgi-bin/getpub?nsf9524.

[27] D. M. Gavrila and L. S. Davis. 3-D model-based tracking of humans in action: a multi-view approach. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition.*, pages 73–80. IEEE Computer Society Press, 1996.

[28] Mark Guzdial and Rick Weingarten. Overview and summary: Research in the union of computer science and education. Workshop report, NSF, 1996. Educational Technology Workshop: Setting a Research Agenda for Computer Science in Educational Technology.

[29] CMU Robotics Institute. Vsam project. www.cs.cmu.edu/∼vsam/.

[30] H. Ishii and B. Ullmer. Tangible bits: Towards seamless interface between people, bits, and atoms. In *Proceedings of CHI 1997*, pages 234–241. ACM Press, 1997.

[31] Prince Kohli, Mustaque Ahamad, and Karsten Schwan. Indigo: User-level support for building distributed shared abstractions. *Concurrency: Practice and Experience*, 8(10), 1996.

[32] Daphne Koller and Mehran Sahami. Toward optimal feature selection. In Lorenza Saitta, editor, *Machine Learning: Proceedings of the Thirteenth International Conference*. Morgan Kaufmann, 1996.

[33] T. K. Landauer, D. Laham, and P. Foltz. Learning human-like knowledge by singular value decomposition: A progress report. In D. S. Touretzky, M. C. Mozer, and M. E. Hasselmo, editors, *Advances in Neural Information Processing Systems 10*. MIT Press, Cambridge, MA, 1998. NIPS97 proceedings, in press.

[34] J. Lee and U. Ramachandran. Architectural Primitives for a Scalable Shared Memory Multiprocessor. In *Proceedings of the 3rd Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 103–114, Hilton Head, South Carolina, July 1991.

[35] Sue Long, Dietmar Aust, Gregory D. Abowd, and Christopher G. Atkeson. Rapid prototyping of mobile context-aware applications: The cyberguide case study. In *Proceedings of the 1996 conference on Human Factors in Computing Systems — CHI'96*, pages 293–294, 1996. Short paper in Conference Companion.

[36] Sue Long, Rob Kooper, Gregory D. Abowd, and Christopher G. Atkeson. Rapid prototyping of mobile context-aware applications: The cyberguide case study. In *Proceedings of the 2nd Annual International Conference on Mobile Computing and Networking*, November 1996.

[37] W.E. Mackay, D.S. Pagani, L. Faber, B. Inwood, P. Launiainen, L. Brenta, and V. Pouzol. Ariel: Augmenting paper engineering drawings. In *Proceedings of the ACM SIGCHI'95 Conference*, 1995.

[38] Jennifer Mankoff and Gregory D. Abowd. Domisilica: Providing ubiquitous access to the home. Technical Report GIT-GVU-97-17, GVU Center, Georgia Institute of Technology, September 1997.

[39] Jennifer Mankoff and Bill Schilit. Supporting knowledge workers beyond the desktop with Palplates. In *Proceedings of the ACM SIGCHI'97 Conference*, 1997.

[40] Jennifer Mankoff, Jonathan Somers, and Gregory D. Abowd. Bringing people and places together. In *Proceedings of the 1998 Spring AAAI Symposium on Intelligent Environments*, 1998.

[41] Jennifer Mankoff, Jonathan Somers, and Gregory D. Abowd. Bringing people and places together with dual augmentation. Paper submitted for review to conference on Cooperative Virtual Environments, 1998.

[42] MIT AI Laboratory. HAL Project. `www.ai.mit.edu/projects/hal`.

[43] MIT AI Laboratory. Intelligent Room Project. `www.ai.mit.edu/projects/hci/hci.html`.

[44] MIT Media Laboratory. Things that Think consortium. `ttt.www.media.mit.edu`.

[45] Michael C. Mozer. The neural network house: An environment that adapts to its inhabitants. In *Proceedings of the American Association for Artificial Intelligence Spring Symposium on Intelligents Environments*, 1998.

[46] Elizabeth D. Mynatt, Maribeth Back, Roy Want, and Ron Frederick. Audio aura: Light-weight audio augmented reality. In *Proceedings of the ACM UIST'97 Symposium*, pages 211–212, 1997.

[47] G. Orwell. *1984*. New American Library Trade, 1989. Reissue Edition.

[48] S. Pakin, M. Lauria, and A. Chien. High Performance Messaging on Workstations: Illinois Fast Messages ( FM) for Myrinet. In *Proceedings of Supercomputing '95*, 1995.

[49] Penio S. Penev and J. J. Atick. Local feature analysis: A general statistical theory for object representation. *Network: Computation in Neural Systems*, 7(3):477–500, 1996.

[50] A. Pentland. Smart rooms. *Scientific American*, 274(4):68–76, April 1996.

[51] A. Pentland, I. Essa, T. Darrell, A. Azarbayejani, and S. Sclaroff. Visually guided animation. In N. Thalmann and D. Thalmann, editors, *Interactive Computer Animation*, pages 143–164. Prentice Hall, 1996.

[52] Umakishore Ramachandran, Mustaque Ahamad, and M. Yousef A. Khalidi. Coherence of distributed shared memory: Unifying synchronization and data transfer. In *18th International Conference on Parallel Processing, St. Charles, Ill*, pages 160–169, August 1989.

[53] Umakishore Ramachandran and M. Yousef Amin Khalidi. An implementation of distributed shared memory. *Software Practice & Experience*, 21(5):443–464, May 1991.

[54] J. M. Rehg, U. Ramachandran, R. H. Halstead, Jr., C. F. Joerg, L. Kontothanassis, and R. S. Nikhil. Space-time memory: A parallel programming abstraction for interactive vision applications. Technical Report CRL 97/2, Digital Equipment Corporation, Cambridge Research Lab, 1997.

[55] T. Richardson, F. Bennett, G. Mapp, and A. Hopper. Teleporting in an X window system environment. *IEEE Personal Communications Magazine*, 1(3):6–12, 3rd Quarter 1994.

[56] Marcel-Catalin Rosu, Karsten Schwan, and Richard Fujimoto. Supporting parallel applications on clusters of workstations: The intelligent network interface approach. In *Sixth IEEE International Symposium on High Performance Distributed Computing (HPDC-6)*. IEEE, Aug. 1997.

[57] H. Rowley, S. Baluja, and T. Kanade. Rotation invariant neural network-based face detection. Technical Report CMU-CS-97-201, CMU Computer Science, 1997.

[58] Gerard Salton. *Automatic Text Processing*. Addison-Wesley, Reading, MA, 1989.

[59] Gerard Salton and Chris Buckley. Term weighting approaches in automatic text retrieval. Technical Report TR87–881, Cornell University, Department of Computer Science, November 1987.

[60] D. Sims. New realities in aircraft design and manufacture. *IEEE Computer Graphics and Applications*, 14(2):91, 1994.

[61] A. Singla, U. Ramachandran, and J. Hodgins. Temporal Notions of Synchronization and Consistency in Beehive. In *Proceedings of the Ninth Annual ACM Symposium on Parallel Algorithms and Architectures*, 1997.

[62] Aman Singla. *Beehive: Application-driven system support for cluster computing*. PhD thesis, College of Computing, Georgia Institute of Technology, 1997.

[63] Quentin Stafford-Fraser and Peter Robinson. Brightboard: A video-augmented environment. In *Proceedings of the ACM SIGCHI'96 Conference*, pages 134–141, 1996.

[64] Roy Want, Andy Hopper, Veronica Falcao, and Jonathan Gibbons. The active badge location system. *ACM Transactions on Information Systems*, 10(1):91–102, January 1992.

[65] Roy Want, Bill Schilit, Norman Adams, Rich Gold, Karin Petersen, John Ellis, David Goldberg, and Mark Weiser. The PARCTAB ubiquitous computing experiment. Technical Report CSL-95-1, Xerox Palo Alto Research Center, March 1995.

[66] Intenational symposium on wearables computers, 1997. `mime1.marc.gatech.edu/wearcon/graphical_wearcon/wearcong.html`.

[67] Mark Weiser. The computer of the 21st century. *Scientific American*, 265(3):66–75, September 1991.

[68] Mark Weiser. Some computer science issues in ubiquitous computing. *Communications of the ACM*, 36(7):75–84, July 1993.

[69] P. Wellner. Interacting with paper on the digital desk. *Communications of the ACM*, 36(7):86–96, 1993.

[70] Andy Wood, Anind Dey, and Gregory D. Abowd. Cyberdesk: Automated integration of desktop and network services. In *Proceedings of the 1997 conference on Human Factors in Computing Systems — CHI'97*, pages 552–553, 1997. Technical note.

[71] Y. Yacoob and L. Davis. Computing spatio-temporal representations of human faces. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 70–75. IEEE Computer Society, 1994.