

MIT.EDU: M-learning Applications for Classroom Settings

Michael Sung, Jonathan Gips, Nathan Eagle, Anmol Madan,

Ron Caneel, Rich DeVaul, Joost Bonsen, Sandy Pentland

Massachusetts Institute of Technology

Media Laboratory

20 Ames Street, Cambridge, MA

{msung, jgips, nathan, anmol, rcaneel, rich, jpbonsen, sandy}@media.mit.edu

Abstract

In this paper we describe Mobile-IT Education (MIT.EDU) classroom applications, which demonstrate the potential of using the MITHril architecture for rapid prototyping of wireless mobile multi-user applications for use in classroom settings. MITHril is a stable, accessible system that combines inexpensive, commodity hardware, a flexible sensor/peripheral interconnection bus, and a powerful, light-weight distributed sensing, classification, and inter-process communications software architecture to facilitate the development of distributed real-time multi-modal and context-aware applications. We demonstrate the power and functionality of this platform by describing a number of MIT.EDU application deployments in educational settings. Initial evaluations of these experiments demonstrate the potential of using the system for real-world interactive m-learning applications.

1. Introduction

Educators and technologists alike are keenly interested in how wireless and mobile technology can enhance the way people learn and interact with each other. It is obvious that these m-learning technologies (e-learning using mobile devices and wireless transmission) can potentially provide important opportunities for learning and collaborative interaction. Reality however, has often failed to live up to these high expectations, and wide technology adoption in classrooms has historically been very slow. Typically, technology is showcased as demonstration systems in various experimental E-classroom initiatives such as [5], but then penetration in general settings is painfully slow. In addition, many m-learning technologies are often limited to content delivery onto mobile devices, missing the rich potential for more interactive learning paradigms.

One cause of limited adoption is that issues such as usability, flexibility, and extensibility are often overshadowed by the need to quickly demonstrate the new features of the technology. The development of the foundation infrastructure necessary to make the technologies most effective is therefore our goal, in order to allow easy deployment of highly interactive and personalized educational technologies.

Several key components are necessary to create such a technology infrastructure for educational settings. First and foremost, a flexible and scaleable system architecture platform is required to be able to appropriately handle classroom settings, potentially involving up to hundreds of individual users. Second, the human factors side of the equation must be properly balanced, and the interface must be appropriately tailored to the application.

MIT.EDU (*Mobile-IT Education*) applications attempt to achieve these goals by building on the MITHril wearable computer technology developed by the MIT Wearable Computing group [6,12] (The name comes from the protective Mithril vests appearing in the Lord of the Rings trilogy). This wearable computing technology supports the rapid prototyping and rollout of large-scale community-based applications, originally with particular application to the support of military and emergency operations. The MITHril project started as an attempt to remedy the substantial human factors, flexibility, and robustness problems plaguing wearable computing research at the end of the 1990s. From these beginnings, MITHril evolved towards a practical, modular system of hardware and software for research in wearable sensing and context-aware interaction [2, 13, 14].

In 2001, the defining feature of MITHril was the modular, distributed, clothing-integrated design based on a unified power/data bus, allowing us to put sensing, computing, and interaction resources where they were most useful and appropriate [11].

In 2003, the advent of inexpensive wireless-capable Linux-PDA hardware allows us to redefine MITHril as a multi-user wireless distributed wearable computing environment, supporting dozens of interacting users and large-scale interaction and sensing experiments. It is this maturation of commercial technology that is paving the road to viable m-learning paradigms.

Digital Anthropology, a cross-registered course between the MIT Media Laboratory and the MIT Sloan Business School, was specifically created in the Spring, 2003 and continued in the Spring, 2004 as a technology tested for investigating deployments of MIT.EDU applications for collaborative learning and teaching feedback.

The syllabus is available through OpenCourseWare, at <http://ocw2.mit.edu/OcwWeb/Media-Arts-and-Sciences/MAS-966Spring2003/CourseHome/index.htm>.



Figure 1: Digital Anthropology seminar with students using the MIT.EDU systems with headsets wirelessly streaming audio data to a remote server, while simultaneously using OpinionMetrics class polling software

Although personal data assistant devices (PDAs) are typically used for individual-user mobile applications, there exists a greater opportunity to tie these PDAs together with a uniform data communication and resource discovery infrastructure that can result in much richer forms of interaction and educational dynamics. The advent of inexpensive wirelessly enabled PDA hardware provides the perfect base platform for multi-user, wireless, distributed wearable computing environment, supporting dozens of interacting users and large-scale interaction and sensing experiments.

We demonstrate the ability to use this standardized PDA hardware tied together with a flexible software architecture and modularized sensing infrastructure to create a system platform where complicated distributed multi-user applications can be developed to enhance educational settings. While our current system implementations are based on PDAs, the software infrastructure is made to be portable to a variety of mobile devices, including cell phones, tablet PCs, and other convergence devices. As such, our systems leverage commercial off-the-shelf components with standardized base-layer communication protocols such as TCP/IP; this allows for the rapid adoption and deployment of our systems into mainstream educational settings.

Even though the variety of applications available for modern mobile devices is quite compelling, they are typically standalone programs with little flexibility or extensibility. We found critical problems in distributed inter-

process communication, signal processing, and sensor data classification that were not addressed by operating systems nor currently available software tools.

The MIThril software architecture addresses these problems by trying to combine the best features and practices from a range of research systems and methodologies, doing so in an open, modular, and flexible way. The three important software systems that form the foundation of the MIThril software architecture are: The Enchantment Whiteboard System for inter-process communication, the Enchantment Signal System for high-bandwidth data streaming, and the Real-Time Context Engine machine learning infrastructure. This software infrastructure ties everything together, allowing network-transparent streaming data communication to arbitrary endpoints capable of real-time, context-aware interaction [2, 12, 13]. These tools address critical needs in the development of mobile applications while imposing minimal constraints on the nature of these applications.

Using the MIThril software, we are able to extract trends and patterns of activity from the environment as well as the individual user. In order to effectively observe contextual data, systems must have a means to gather, process, and interpret this real-time contextual data. To facilitate this, the MIThril architecture includes modular sensor hubs that can be used to instrument devices for contextual data gathering.

MIT.EDU applications are layered upon this base capability, providing a way to gather these heterogeneous streams of information, perform real-time processing and data mining on this information, and return classification results and statistics that can be used for educational applications. This important contextual data can then be used to derive meaningful information such as immediate user context/profiles, user interaction patterns, aggregated user statistics, and even social network topology or organizational structure. This information can result in more effective, context-aware applications that can augment classroom dynamics in educational settings.

In this paper, we first describe the hardware and software infrastructure of the MIThril architecture in Section 2. After this brief technical introduction, we describe a number of case studies of MIT.EDU applications in Section 3, to demonstrate the power and flexibility of the architecture. Our conclusions can be found in Section 4.

2. MIThril Technical Overview

MIT.EDU applications use the MIThril architecture to provide distributed applications in classroom and collaborative settings. Here we provide a basic description of the hardware and software architecture; please refer to DeVaul et al. [6] for a more detailed technical description.

2.1 Hardware

The MITHril hardware architecture is a highly flexible, modular system that is tied together by wired/wireless networking protocols and a unified multi-wired protocol power/data bus for sensors and peripherals. An example of the MITHril configuration is laid out in Figure 2.

MITHril employs the Sharp Zaurus SL-5500 PDA for applications requiring real-time data analysis, peer-to-peer wireless networking, full-duplex audio, and graphical interaction. The Zaurus SL-5500 is a complete embedded Linux system. It features a 206-Mhz StrongARM processor, 64 MB SDRAM, 16 MB Flash, CF and SD expansion slots, full duplex audio, qVGA color touch screen, and an integrated QWERTY keyboard. The StrongARM provides sufficient computing power for moderate bandwidth signal processing and real-time classification.

The MITHril system supports wireless networking through the Zaurus CF interface using the 802.11b wireless protocol. This low-cost wireless networking capability is a crucial enabling feature, allowing us to implement multi-node, distributed wearable applications. In general, the CF card slot allows for a rich variety of peripherals/sensors, including cell-phone modems, image and video cameras, Bluetooth and 802.11b (WiFi) wireless, and even head mounted displays (HMD). The Zaurus also provides an RS232 serial port, which we use to interface with the SAK2 sensor hub to communicate with sensors.




Figure 2: MITHril system, composed of the Zaurus PDA (right) with sensor hub and physiological sensing board (top), EKG/EMG/GSR/ temperature electrodes and sensors (left) and combined three-axis accelerometer and IR tag reader (bottom)

The SAK2 sensor hub is responsible as bridge to the sensor data, providing sensor data acquisition, buffering, and sequencing. The SAK2 provides data acquisition capabilities in the form of a CF storage card interface, a real-time clock for time-stamping, flexible battery power regulation, and 2.4 GHz wireless transceiver. The SAK2 can also be used without the Zaurus as a stand-alone data acquisition system. This is particularly useful for large-

group applications that do not require real-time processing, WiFi wireless or complex user interaction.

MITHril leverages a wide variety of sensors, including accelerometers (for motion detection), IR active tag readers (for location and ID), battery monitors, GPS units, audio microphones, and physiological sensing including EKG/EMG, galvanic skin response (GSR), temperature, heat flux, pulse oximeters, and polar heart monitors.

2.2 Software

The MITHril software architecture allows designers to quickly prototype distributed, group-based applications that use contextual information about the members of a group. Layered on top of standard libraries, this middleware comprises three important parts: the Enchantment Whiteboard, the Enchantment Signal system, and the MITHril Real-Time Context Engine. Respectively, these three layers provide the ability to easily coordinate between distributed applications, transmit high bandwidth signals between applications, and create classification modules that make a group's changing contextual information available to applications.

The Enchantment Whiteboard system is a distributed, client/server, inter-process communication system that provides a lightweight way for applications to communicate. In a similar fashion to the Linda [3], it processes, publishes, and receives updates from tuple spaces where information is decoupled from specific processes. This is particularly useful in mobile, group based applications where group members may not be known a priori and may come and go over time.

For higher bandwidth signals, especially those related to the sharing and processing of sensor data for context aware applications, we developed the Enchantment Signal system. The Signal system is intended to facilitate the efficient distribution and processing of digital signals in a network-transparent manner. The Signal system is based on point-to-point communications between clients, with signal "handles" being posted on Whiteboards to facilitate discovery and connection. In the spirit of Whiteboard interactions, the Signal API abstracts away any need for signal produces to know who, how many, or even if, there are any connected signal consumers.

The MITHril Real-Time Context Engine [7] is an open-source, lightweight, and modular architecture for the development and implementation of real-time context classifiers for wearable applications. Using the context engine, we can implement lightweight machine learning algorithms (capable of running on an embedded system like the Zaurus) to process streaming sensor data, allowing the MITHril systems to classify and identify various user-state context in real-time.

3. MIT.EDU Applications

The MIT Wearables Group has about forty MIThril systems in active use, including group-based applications within several collaborative class settings involving thirty simultaneous users. The authors developed an ongoing seminar course called Digital Anthropology in order to develop a tested for deploying prototype MIT.EDU applications as well as to facilitate interaction with the MIT Sloan Business School for larger real-world class deployments.

In the following section, a number of real-world case examples of multi-user MIT.EDU applications that are built upon various parts of the MIThril platform are detailed. We first start with a description and evaluation of OpinionMetrics, a distributed application that demonstrates the power of real-time student and teacher feedback. This is followed by brief descriptions of our Reality-Mining, Socio-PhysioMetrics, and GroupMedia projects, examples that were chosen to showcase a few novel applications with unique real-time interaction, analysis, and feedback possibilities that illuminate the potential of m-learning technologies. These examples demonstrate the modular, configurable nature of the MIThril hardware and the flexibility of the software architecture to accommodate a variety of high bandwidth, real-time applications.

3.1 OpinionMetrics

In an MIT classroom, students attend an introductory Finance lecture. In front of each student, a WiFi equipped Sharp Zaurus displays a simple, functional interface with buttons labeled “Applause”, “Bored”, and “Lost”. During the lecture, the students use this interface to provide instant feedback on their status in class. The lecturer is also equipped with a Zaurus that allows her to monitor this feedback in real-time. When the number of “Lost” students reaches a predefined threshold, the lecturer’s Zaurus flashes red as a warning that the material isn’t being absorbed.

To clarify the class’ level of understanding, the lecturer asks a question. Her Zaurus-equipped teaching assistant polls the class with the push of a button. When he starts the poll, the students’ interface raises a poll dialog box. When the poll ends, the results are displayed on everyone’s Zaurus, and the lecturer has a better idea of the class’ understanding. This information is logged for later analysis, allowing the lecturer to craft better presentations.

OpinionMetrics is a set of applications designed to provide lecturers and teaching assistants with feedback about how students are tracking material presented in class. The system has been tested in multiple classroom

situations, and the feedback for the potential to improve the learning environment is positive.

We have successfully deployed OpinionMetrics in several trial runs in actual classroom settings at MIT. One particular system rollout was during a neurobiology class with over twenty students, where each student’s OpinionMetrics data as well as the class’ audio were recorded. The OpinionMetrics systems scaled very well under the resource limitations of the available 802.11 infrastructure. Despite the fact that every student already had a WiFi-enabled Tablet PC, there were no problems running so many simultaneous systems streaming OpinionMetrics data.

The data logged from the class from each individual is plotted in Figure 3, and the aggregate results of this data are displayed in Figure 4.




Figure 3. (Top) Individualized OpinionMetrics button press data collected from eighteen students in a lecture.

We can make several interesting observations from the data collected. From the individual data, we can determine which students felt lost or interested, and the times when they felt compelled to indicate their class state. These timestamps can be synchronized to points in the recorded audio stream to identify the topics that the students struggled on or found interesting. This information can allow a teacher to tailor the educational content and resources to better fit the needs of particular students.

From the aggregate data, we can see spikes of interest and confusion during specific points in the lecture. This real-time information is potentially very useful to determine the average state of the class. The OpinionMetrics software was set to notify the lecturer once a critical simultaneous “lost” threshold was reached. In the class, the lecturer actually responded to these warnings by issuing

in-class polling questions, which were sent directly to the student systems. The lecturer can use the results of these polls to further identify the class state and potentially change his lecture material dynamically.

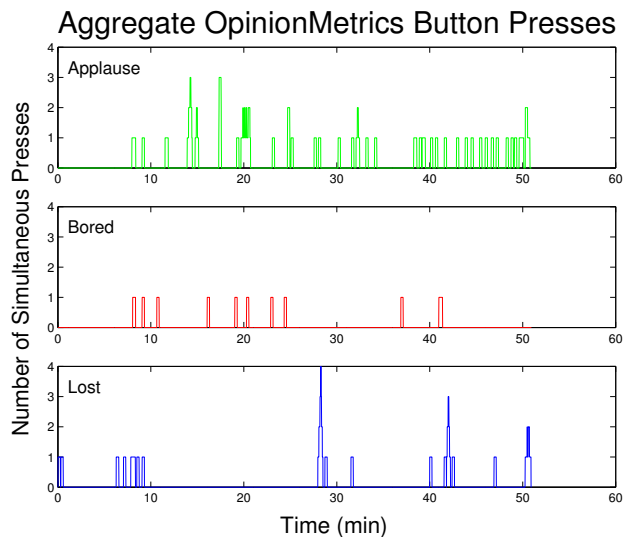


Figure 4: The same data is displayed as three aggregates where each button press is counted for a fixed number of seconds (fifteen in this case).

A student survey of the usability of the OpinionMetrics systems was conducted. The survey showed that the students generally responded very favorably to the anonymous feedback interaction that OpinionMetrics offered. They felt that they had a significantly better way to express their understanding of the material and found that the interaction level between the students and professor was raised. In general, they did not feel that having the OpinionMetrics system in front of them created a major distraction and did not find it uncomfortable to hold or use. On average, students wanted to use the system in other classes as well. The lecturer of the class was also very interested in the potential of the MIT.EDU systems. One suggestion that was made was to add alternative output modalities such as audible or tactile cues (instead of strict visual graphs) in order more naturally gauge class state without distraction. The summary of the survey responses is presented in Figure 5.

There are other projects that have goals similar to OpinionMetrics that allow students to give feedback and allow a lecturer to take polls [1, 9]. While some of these systems are relatively low-cost and allow for large-scale implementations, OpinionMetrics has an advantage because of its customizability and ability to run on devices that students may already own such as laptops and PDAs.

We have received positive feedback from the OpinionMetrics project, and there is clearly a use for such a

system in lecture halls where individual expression can be hindered by large numbers of people competing for attention and the fear of exposing one’s lack of understanding to peers.

It was <i>rather</i> easy to hold the device.	5.5 +/- 1.4
It was a <i>minor</i> distraction during lecture.	3.8 +/- 1.6
The bored button was used <i>rarely</i> .	1.6 +/- 1.1
The interested button was used <i>occasionally</i> .	3.1 +/- 2.2
The lost button was used rarely.	2.2 +/- 1.8
The ability to give anonymous feedback is <i>extremely</i> helpful.	6.1 +/- 1.2
The system increased students' ability to express feelings about lecture at any given moment by a <i>great deal</i> .	5.9 +/- 1
The system increased student-teacher interaction <i>significantly</i> .	5.3 +/- 1.4
The system made class time more effective by a <i>bit</i> .	4.3 +/- 1.5
The system increased learning by a <i>bit</i> .	3.9 +/- 1.8
Students <i>really</i> want to use the system in other classes.	5.7 +/- 1.3

Figure 5: Summary of survey results given to an eighteen-person class after using OpinionMetrics (out of a 7 point scale). The words in *italics* are our interpretation of the students’ average rating, which is given in the second column along with the standard deviations.

3.2 Reality Mining and Conversation Analysis

Throughout the semester, every student in the Digital Anthropology course was outfitted with a MIThril system, which recorded continuous high quality streaming audio using the Enchantment infrastructure to a remote server. Profiles of a participant’s conversation such as speaking rate, energy, duration, participants, interruptions, transition probabilities, and time spend holding the floor were calculated using conversation detection and analysis algorithms. This information gives valuable insight into the context and content of conversation as well as capturing the dynamics of how such conversations are structured.

Another interesting area of exploration is deploying the systems in group settings to capture conversations and develop statistics on the dynamics of group interaction. Classroom dynamics play a crucial role in the success of a class. It is common knowledge that an instructor must have a sense of how many people are following the lec-

ture, who could use more personalized help, and which parts of lecture need to be reviewed. In an effort to substantially augment an instructor’s intuition, our system can quantify these types of metrics to a much greater level of precision.

Each standard MIThril system was augmented with an MIT.EDU application to record audio continuously, storing it locally until it could be transmitted to a server over an available 802.11b network. Participants continuously provide subjective feedback on comments and discussion using a 2D touch pad. The feedback interface converts the task of providing continuous feedback into a low-attention, secondary task.




Figure 6: A one-hour class with voicing segments mapped above aggregate interest level

By correlating peaks in interest/approval with the individual audio inputs, the system can automatically provide a summary audio track consisting of comments that had high approval or interest ratings, and employ speech analysis to identify topics that had high (or low) ratings. Dynamic maps of student interaction can be generated and publicly displayed to reflect the roles and dyadic relationships within a class. This analysis can help develop deeper insight into the underlying dynamics of the class.

Once detected, the conversation audio streams are extracted and analyzed. Figure 7 shows a selection of features that can be gleaned from this audio data. Profiles of a student’s typical classroom behavior are built over time using conversation features such as speaking rate, energy, duration, participants, interruptions, transition probabilities, and time spent holding the floor. By comparing relative volume levels of a student’s voice in multiple microphones, it even becomes possible to infer physical proximity to an approximate degree [10].

This system uncovers information concerning the effectiveness of the class, as well as the dyadic relationships between individuals. The information collected includes a list of the peers that a student typically sits by, avoids, talks to, interrupts, and transitions. As can be seen from Figure 8, a professor (s9) is obviously the dominant member while his advisees (s2, s7, s8) concede the floor

to him with relatively high probability - indicative of his influence.

Speaker Number	Floor time (%)	Avg Comment (sec)	Nearest Neighbor	Transition (Name, %)	Avg Inter-est	Group Interest
s1	1.5	4.1	s8	s8-27	.21	.44
s2	2.2	2.2	s9	s9-47	.13	.36
s3	9.9	3.5	s9	s4-22	.20	.22
s4	11.4	9.6	s7	s6-23	.05	.30
s5	12.8	8.8	s7	s9-37	.18	.33
s6	16.9	6.6	s4	s7-28	.09	.21
s7	10.1	6.6	s4	s9-30	.19	.24
s8	10.8	10.9	s1	s9-26	.40	.32
s9	24.4	6.9	s7	s6-22	.17	.25

Figure 7: Metrics for classroom interaction analysis

Subjective feedback is pooled and shared with the participants via a public display. Comments that give rise to wide variations in opinion cause the discussion to focus on the reason for disparate opinions, and controversial topics can be retrieved for further analysis and debate. Opinions and comments can also be clustered using ‘collaborative filtering’ to display groupings of opinion, allowing within-group and between-group debate. This project demonstrates our ability to capture extremely rich data on everyday behavior within the classroom.

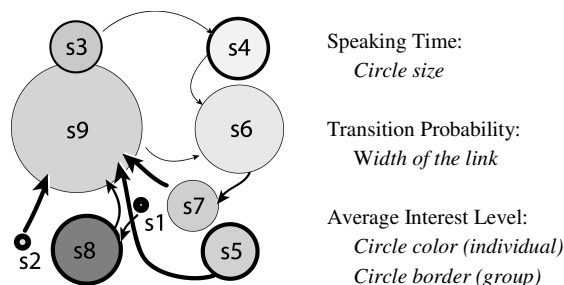


Figure 8: A visual depiction of the professor (s9) and student dynamics

We have also deployed our systems in MIT Sloan Business School’s negotiation classes in the Fall, 2003, where it is useful to be able to monitor individual and group reactions to structured interactions to analyze conversational dynamics using our techniques as well.

We hope to use the information obtained from these controlled experiments to measure the extent to which it can be leveraged to create a more effective and informative classroom experience.

3.3 Socio-PhysioMetrics

The Digital Anthropology seminar is hosting a special guest speaker. The speaker and audience members are wearing physiological monitoring systems that measure skin temperature, EKG, and galvanic skin response. Behind the speaker, a projector displays these aggregated physiological signals in real-time, allowing the audience and speaker to gauge the effects of the speech on the audience as well as providing biofeedback to the speaker.

The guest speaker and the members of the audience are also wearing headset microphones and WiFi-enabled Zaurus that stream the individual audio signals to a central server for logging, playback and later analysis. The separate audio signals for each speaker make analyzing the structure of conversations easy. Some members of the audience are wearing an experimental motion-based head nodding/shaking measurement system, which can classify whether a person is subconsciously in agreement/disagreement in real-time. This data is also being displayed on the projector and recorded for later interpretation.

For our Digital Anthropology classes, the speaker and audience members were outfitted with MITHril systems that can measure skin temperature, EKG, and galvanic skin response. The MIT.EDU applications used these measurements, plus an accelerometer-based real-time head nodding/shaking classification systems developed using the MITHril Context Engine framework. These systems allowed us to stream physiologic/movement data in conjunction with the opinion data from the OpinionMetrics software. Figure 9 shows one of the authors wearing a MITHril system with MIT.EDU application that streaming real-time data to a nearby projector.

The data can be used to gauge interest and agreement levels in real-time, and to be able to cross-compare self-reporting results to baseline information such as unconscious nodding in agreement and psycho-physiologic cues such as heart rate and galvanic skin response, which is highly correlated with stress and sympathetic nervous system arousal.

The real-time visualization of the aggregated Enchantment signal information serves a dual-roll for the speakers and teacher. The audience's aggregated psycho-physiology statistics can be used to gauge audience attention and interest in the form of socio-biofeedback. The speaker can also be wired with his own physiological signals, which is useful as biofeedback for the speaker to identify his own mental state, or by the teacher, who can observe the effects of the speaking or negotiation interaction. Within the context of a structured interaction of a speaking or negotiation class, the person's performance can be studied in a very controlled environment, giving the individuals as well as the instructor valuable feedback.



Figure 9: MIT.EDU application performing real-time FFT analysis and classification on accelerometer data, and wirelessly streaming real-time EKG/GSR/ temperature and classification results to a remote projector

These socio-physiometric displays have the power to radically change the dynamics and interaction of students and teachers, as well as to provide additional dimensions of information dissemination. Similar displays that show real-time group contextual information such as speaking time of participants in a meeting have been shown to have positive impacts on the dynamics of group interaction [8].

This project uses the Enchantment Signal system in addition to the Whiteboard system, to provide an intuitive infrastructure for producing and consuming heterogeneous streams of information. The Signal system, in conjunction with the Enchantment Whiteboard, allows us to simply plug in a variety of sensors, and have those streams of information simply flow from producer to consumer in a network-transparent way. In addition, very high data bandwidths are achievable (allowing for applications such as VoIP-style communications), and can easily scale up to the bandwidth limits of the underlying wireless 802.11 network.

3.3 GroupMedia and Social Influence/Interest

Group deliberations and decision-making are an integral aspect of Sloan Business School. Four business school students are keenly involved in an animated discussion to find a class project they have to execute as a team. Individuals are using Rateit! on their Zaurus PDAs to give an objective rating of how interesting they find the ideas. This can be correlated with their head-movement and nodding, speech features and physiology, to understand the socio-physiometrics of brainstorming and idea-generation.

The concepts of conversational analysis and socio-physiometrics can be combined to analyze the interaction

of groups, which we dub GroupMedia. The GroupMedia system has been used to measure conversational interest in ten sessions of our Digital Anthropology class. Each session involved a group of three or four people, for durations of between ten minutes and one hour. The students engaged in conversation and brainstorming sessions, while we recorded their audio dynamics, head movement, and subjective interest ratings.

Head movement and an individual's speech features were easily obtained through motion sensors and microphones that plugged into our MIT.EDU systems. Subjective interest rating was obtained through an interactive MIT.EDU application called *Rateit!* This application is a variant of the OpinionMetrics software, but with a particularly designed GUI so that the users can change interest ratings without significant cognitive load and without having looking at the PDA touch-screen to use it.

We can analyze the voice streams of the individuals of the group session to determine the 'influence' that individuals have on each other. This influence parameter expresses then how strong the overall state for a person is depending on the state of another person. In this case we use a simple two-state model of speaking vs. not-speaking dynamics of the recorded audio to model individual dynamics, and then measure the influence parameter to determine the coupling between speakers. In Choudhury and Pentland [4] this measure of influence was shown to have an extremely high correlation with one measure of the social novelty (and thus presumably the interestingness) of the information being presented.




Figure 10: Plot of participant interest ratings and the corresponding influence parameters.

From the derived influence parameters and the objective interest ratings, we can begin to see these trends relating social influence and individual interest. Figure 10 shows a graph of the typical data from these meetings. The graphs show group interest ratings during an eight-minute session, and the corresponding influence parameters calculated from the audio features. A long-term rise in interest is observed, along with a corresponding long-term rise in influence parameters; the influence parameters are a statistically significant predictor of user interest ratings with a correlation of $r=0.5$. More importantly, the

'bumps' in the interest rating correspond to 'bumps' in the influence parameters. This supports the ideas that as individuals find the conversation more interesting they begin to influence each other more (were more 'engaged').

There was also a correlation between the overall head-movement in the group and objective interest. Figure 11 shows clusters of head nodding behavior when there are changes in interest ratings. Bursts of group head-nodding correctly identify 80% of the changes in group interest level, with a 1/3 false alarm rate. Head nodding is not a perfect indicator of interest change, nor does it give the sign of the change, but it does provide the very useful information that we should look for changes in participant behavior and interest.

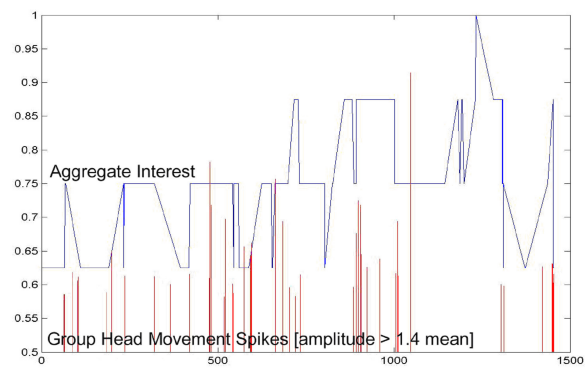


Figure 11: Plot of interest ratings and the corresponding overall group head-nodding energy

4. Conclusions

The MIT.EDU applications we have described demonstrate the potential of a flexible system infrastructure capable of a variety of individual and group-based context-aware applications. As the various applications that have already been developed demonstrate, the MITHril hardware and software infrastructure allow for the rapid implementation of complex, distributed applications. The fact that the infrastructure allows the rapid prototyping of new MIT.EDU applications with minimal effort means that we can conceive and iterate on new applications very easily, further simplifying the design process.

The responses we have had to the MIT.EDU applications provide evidence that we can create significant positive changes in the educational dynamics of our classrooms. Of special interest are the real-time and post-processed statistics that can be gathered with our infrastructure that can help teachers to more effectively guide and direct their teaching interactions with students. These statistics can reveal important trends in classroom behavior, and result in more efficient dynamics where more effective learning can take place.

It is our hope that other people will find our infrastructure and software tools useful. As with all MIThril code, these tools are open-source and available under the terms of the GPL, along with our hardware design files for the SAK2 sensor hub and other sensor designs [12].

5. Acknowledgements

We would like to recognize the Media Laboratory wearable community for their contributions to the MIT.EDU applications. A special thanks goes to Ivan Chardin, who helped with the MIT.EDU system integration and applications, and Joost Bonsen, who originated and the Digital Anthropology seminar as a way to test and develop MIT.EDU applications.

6. References

- [1] J. Boyle, et. al., "Experience with Classroom Feedback Systems to Enable Socratic Dialogue in Large Engineering Classes", *Engineering Education 2002, Professional Engineering Scenarios, IEE*, 2002
- [2] P. J. Brown, J. D. Bovey, and X. Chen, "Context-Aware Applications: From the Laboratory to the Marketplace", *IEEE Personal Communications*, October 1997
- [3] N. Carriero and D. Gelernter. "Linda in Context", *Comm. ACM*, April 1989
- [4] T. Choudhury, "Sensing and Modeling of Human Networks", Ph.D. Thesis, Department of Media Arts and Sciences, MIT, 2003
- [5] S. M. Davis, "Research to Industry: Four Years of Observations in Classrooms Using a Network of Handheld Devices," *Wireless and Mobile Technologies in Education*, 2002.
- [6] R. W. DeVaul, M. Sung, J. Gips, and A. S. Pentland, "MIThril 2003: Applications and Architecture", *International Symposium of Wearable Computers*, 2003
- [7] R. W. DeVaul and A. S. Pentland. "The MIThril Real-Time Context Engine and Activity Classification", *Technical Report, MIT Media Lab*, 2003.
- [8] J. M. DiMicco, W. Bender, "Second Messenger: Increasing the Visibility of Minority Viewpoints with a Face-to-face Collaboration Tool", *Intelligent User Interfaces 2004*
- [9] R. Dufresne, et. al., "Classtalk: A Classroom Communications System for Active Learning", *Journal of Computing in Higher Education*, 1996
- [10] N. Eagle, A. Pentland, "Social Network Computing", *The Fifth International Conference on Ubiquitous Computing (UbiComp)*, 2003.
- [11] P. Lukowicz, U. Anliker, G. Troster, S. J. Schwartz, and R.W. DeVaul, "The WearARM Modular, Low Power Computing Core", *IEEE Micro*, 21(3):16-28, May-June 2001
- [12] MIT Wearables Lab, www.media.mit.edu/wearables
- [13] J. Pascoe. "Adding Generic Contextual Capabilities to Wearable Computers", *Digest of Papers. 2nd International Symposium on Wearable Computers, IEEE*, October 1998
- [14] T. Starner, B. Schiele, and A. Pentland, "Visual Context Awareness in Wearable Computing", *In Digest of Papers. Second International Symposium on Wearable Computers*, pages 50-57. *IEEE Computer Society*, October 1998