

ORNET: A Network for the Operating Room of the Future MIT2000-10

Progress Report: January 1, 2002—June 30, 2002

John Guttag and Hari Balakrishnan

Project Overview

The overall goal of this project is to start to bring the benefits of ubiquitous computing and communications to health care in hospitals, physician's offices, homes, and the field.

Computing already pervades medicine. Processors can be found in electrocardiogram machines, blood pressure monitors, drug pumps, and even stethoscopes. Each device greatly facilitates medicine's fundamental process of gathering information about the state of a patient and then using that information to choose interventions intended to improve the patient's state.

Unfortunately, all of this enabling technology has combined to create a disappointing system in which a large number of excellent but isolated point solutions are often considerably less than the sum of the parts. The problem starts with the architecture of the typical medical device:

- A transducer (the probe) that measures an analog signal (e.g., electrical voltage or reflected sound) emanating from the patient,
- A wire connecting the transducer to,
- Special purpose hardware,
- A program that runs on the special purpose hardware and processes the data gathered by the transducer, and
- A dedicated user interface.

This leads to a number of limitations, including the inability to easily fuse data, share resources, upgrade algorithms, and systematically collect data for further analysis. In short, the devices don't function as a system.

A system links various components together to achieve higher-level goals. The utility of the system, in any given situation, is determined by how well the different components inter-operate to perform the desired tasks. Generally, an efficient system should facilitate and control information flow between components while minimizing any unnecessary replication of resources or functionality. An opportunity exists to re-architect today's inefficient medical system by leveraging new technologies and applying this classical systems viewpoint.

One factor, as mentioned above, is the ubiquity of computing and communication. Each will soon be everywhere. Taking advantage of this opportunity requires a multi-level architecture that makes appropriate use of a variety of communication and computation technologies.

A second factor is the ever-increasing selection of smaller less-invasive sensors that permit extensive gathering of physiological data. Taking advantage of this opportunity requires an architecture that scales in the number of sensors. Clearly, connecting each sensor by a wire to dedicated user interface is not such an architecture. Furthermore, the architecture must be plug-and-play.

A third factor is the availability of increasingly sophisticated algorithms for analyzing signals and data. Taking advantage of this opportunity requires an architecture that provides a clean separation between the parts of the

system that gather the initial data (i.e., the probes) and the parts of the system that process the data in real time and save the data for subsequent analysis.

A fourth factor is the rapid improvement in the capabilities of commodity hardware for both computation and storage. Taking advantage of this opportunity requires moving away from the special purpose hardware commonly used in today's medical instruments.

This project involves building new system for medical networks of devices that supports a variety of devices in a reliable, yet flexible manner. As part of this project we are developing techniques to facilitate data fusion and increase the reliability of proxy functions and developing protocols for device registration that yield plug-and-play functionality. Finally, we are working with clinicians to build interesting and useful applications of this technology.

Progress Through June 2002

This project evolved from the SpectrumWare project which was funded, in part, by NTT. In SpectrumWare we pushed the boundary between software and hardware in the communications domain. Here we are doing much the same thing in a different domain. A medical device consists of a sensor or actuator connected to an A/D converter and a (usually) low cost low power wireless network interface.

In the first phase of this project, we designed a multi-tiered architecture. Our architecture employs a collection of gateways and software proxies to bring sensors and actuators onto the network. All information processing occurs in the proxies, which are chained together to perform data fusion and create new virtual devices. User interfaces are completely separated from both devices and proxies. This facilitates remote monitoring using user interface devices ranging from large format high resolution displays to 3G telephones.

In the next phase we gathered some experience with all aspects of the system, by building a relatively primitive, but comprehensive, prototype. This included

- A prototype device and network manager. This software is responsible for managing the network. This includes detecting the arrival and departure of devices, connecting proxies to devices and to each other, supplying data to remote monitors, and providing security. The relatively small scale of the network makes it possible to centralize much of the control.
- Communication protocols,
- Several physical devices, including an invasive blood pressure device, a network controlled drug infusion pump and an electronic stethoscope.
- Device proxies for the physical devices and for a number of virtual devices, e.g., a cardiac monitor driven by the Massachusetts General Hospital cardiac data base,
- Synthetic devices, e.g., a robust heart rate monitor, built by connecting proxies for multiple physical devices, and
- Wireless remote monitoring communications and application software.

The experience gathered during this phase helped us to solidify our overall system architecture. We are now confident that we have a framework that will serve us well going forward.

In the most recent phase of the project we made progress in the following areas:

- Improved our prototype system
- Built a sensor that combines an electronic stethoscope with an ECG and connected it wirelessly to our network
- Designed patient-specific algorithms to detect the onset of epileptic seizures
- Designed algorithms to analyze cardiac abnormalities
- Began work on an alerting system

In order to evaluate our system in more real-world conditions, we are planning to develop proxies to assist clinicians in analyzing EEG data in real-time. This is important in localizing the focus of an epileptic seizure. To this end we looked at EEG signals collected and classified by clinicians at Children's Hospital and we

implemented several different algorithms for detecting the onset of epileptic seizures. Currently we are testing the algorithms against a database of seizures. These algorithms will be used in proxies to analyze data collected from patients in real-time. When the onset of a seizure is detected our software will notify clinicians and inject radioactive dyes so that the focus of the seizure can be found via SPECT.

Using data collected and classified by clinicians at Massachusetts General Hospital, we are developing algorithms to classify heart murmurs. Murmurs can indicate serious conditions, such as mitral valve prolapse, or less serious conditions, such as aortic stenosis. Current practice requires a relatively costly procedure, echo cardiogram, to distinguish among these murmurs. We are seeking to develop software to distinguish the various classes of murmurs, based on easily collected acoustic and electrical data.

Work on our prototype system included a redesign to make it more robust and the integration of a new acoustic sensor. This sensor is more sensitive than other electronic stethoscopes. We have integrated this stethoscope with a simple ECG so that we can easily map the acoustic signal to parts of the heart cycle. The time correlated data streams from these sensors are sent over a wireless network connection to software proxies.

In order to assure properties related to safety criticality, we started work on a principled and systematic way to build alerting systems. Our software architecture links sensors and actuators with patients and caretakers and provides sophisticated mechanisms for detecting problems and raising alerts. An ideal alerting system would detect every problem, issue no false alarms, provide sufficient advanced notification to allow for preventive interventions, and warn medical personnel when an intervention has the potential to preclude potentially useful future treatments.

A perfect alerting system is not possible. There is always uncertainty about the quality of data, the interpretation of the data, and future events. The detrimental effects of this uncertainty clearly manifest themselves in existing medical alert systems. Most noticeably, false alarms are rampant, as any visit to an ICU makes clear. Missed detections also occur, sometimes with catastrophic results.

We are starting by designing a formal model of the alerting system for a number of medical environments. It is based on an approach devised by Jim Kuchar of MIT's Astro and Aero Department. It converts the problem to the domain of signal-detection theory, where one must determine if a known signal is present in background noise. He models uncertainties using probability density functions, from which the probabilities of false detections and missed detections can be calculated. This mechanism enables informed decisions on how to set appropriate policies.

Research Plan for the Next Six Months

We do not anticipate receiving funding from NTT for this project during this period.