

ORNET: A Network of the Operating Room of the Future

MIT2000-10

Progress Report: July 1, 2001—December 31, 2001

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. We plan to start with the operating room, but expect to quickly extend our work to other venues ranging from doctors offices, to homes, to the field.

With upwards of fifty microprocessors in every hospital operating room, 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 chose 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 a 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 give 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 realtime 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 December 2001

This project evolved from the SpectrumWare project. 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 on devices ranging from large format high resolution displays to 3G telephones

We next tried to gather some experience with all aspects of the system, by building a relatively primitive, but comprehensive, prototype. This included

- A prototype device 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.
- 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.

At the end of the year, we produced a five minute video that describes the overall project. This video is available with Japanese narration.

Research Plan for the Next Six Months

Our first activity will be to document the work of the previous six months in one or more technical papers or reports. We will then move on to solidify our infrastructure software, work with clinicians to gather data, build novel proxies that provide new functionality, and port our remote monitoring capability to 3G telephones.

Our prototype software was not built with either scalability or reliability in mind. Over the next year we plan to invest considerable effort in rectifying this. We will use redundant hardware and a fail over model to provide fault tolerance for hardware faults. Static and dynamic checking will both be used to reduce software faults.

We have already started talking with clinicians about gathering data. Over the first four months of the year we expect to gather acoustical and EKG data from adults with cardiac anomalies and EEG data from children who suffer from seizures. Working from this data we hope to develop novel proxies that will provide allow clinicians to use simple instruments to provide unprecedented information. Initially we plan to focus on diagnosis of mitral valve abnormalities using stethoscopes and early detection of seizures using scalp EEG.

Our LAN-based remote user interfaces running on a hand held computer have been well received by the clinicians to whom we have shown it. However, while this will be extremely useful within hospitals or medical offices, physicians would like to be able to easily monitor patients from locations where there is no LAN. We therefore plan to provide an interface based on the telephone network.