

TECHNICAL PROPOSAL

Table of Contents

1. TECHNICAL DISCUSSIONS	2
A. Statement of Work	2
A.1. Abstract	2
A.2. Objectives	3
A.3. Approach	3
A.4. Methods	4
A.4.1 Base Technologies	4
A.4.2 Components	9
A.4.3 Integrating the Components	10
A.4.4 Expanding the Scope	13
A.4.5 Energy Management	14
Related Work	14
References	15
A.5. Schedule	Error! Bookmark not defined.
B. Personnel	Error! Bookmark not defined.
B.1. Principal Investigator	Error! Bookmark not defined.
B.2. Additional Investigators	Error! Bookmark not defined.
B.3. Faculty	Error! Bookmark not defined.
B.4. Research Staff	Error! Bookmark not defined.
B.5. Research Assistants	Error! Bookmark not defined.
B.6. Resumes	Error! Bookmark not defined.
2. OTHER CONSIDERATIONS	<i>ERROR! BOOKMARK NOT DEFINED.</i>
3. INFORMATION TECHNOLOGY SYSTEMS SECURITY	<i>ERROR! BOOKMARK NOT DEFINE</i>

1. TECHNICAL DISCUSSIONS

A. Statement of Work

A.1. Abstract

In response to the National Library of Medicine's Broad Area Announcement titled "Application of Advanced Network Infrastructure Technology in Health and Disaster Management," the Decision Systems Group (DSG) at Brigham and Women's Hospital has proposed the development of SMART, acronym for Scalable Medical Alert and Response Technology, as a model for both local area and wide area patient monitoring. It will also serve as a model for patient monitoring in a disaster situation. To test this model, the DSG has proposed building a test bed patient monitoring system in the Emergency Department at the Brigham and Women's Hospital.

In support of the Brigham and Women's Hospital SMART proposal, as a subcontractor, MIT's Laboratory for Computer Science (LCS) will design a scalable location-aware patient monitoring system. The three main components of this system are (1) a patient-based sensor system; (2) an indoor location infrastructure for tracking equipment and people; and (3) a location database. This patient monitoring system will be designed in a way that ensures it will work seamlessly across indoor and outdoor contexts.

This system will be based on several research initiatives of the Networks and Mobile Systems Group within LCS. The patient-based sensor system will be based on the ongoing Patient Centric Network project [Harfst 02]. The Cricket system [Priyantha 00] will be the base for the indoor location and tracking system. It will be extended with commercially available RFID technology for equipment tracking. INS [Adjie-Winoto 99] will be used for the location database. The effort will be integrated with SLAM, a new initiative examining scalability issues in sensor networks and resource location and tracking.

The Patient Centric Network is an ongoing research project: its architectural concepts and software are ready for testing outside the laboratory. The Cricket and INS systems are established research projects ready for initial deployment in production contexts. RFID is a commercially available tracking technology. The SLAM project will provide scalable data management and query processing techniques as it matures.

LCS will work closely with researchers and practitioners at BWH and CIMIT to integrate the location-aware patient monitoring system with a decision support system that will be developed by BWH researchers to recommend proper allocation of personnel and material resources for patients with cardiovascular and respiratory complaints.

A.2. Objectives

Today there are many challenges in delivering health care including

- Scarcity of trained personnel,
- Overworked providers,
- Unpredictable numbers of patients,
- Pressures to contain costs, and
- Demographic shifts: an aging population and their healthcare issues

These challenges lead to the desire to develop computer and networking systems to augment capabilities of the available healthcare workforce to serve the increasing number of patients. To alleviate some of these pressures, we are proposing to develop a system that will help to provide a higher standard of care, by

- Continuously monitoring the physiological status of at risk patients,
- Providing a ubiquitous communications link between the patient monitors and health care professionals,
- Tracking the location of patients and health care resources,
- Automatically alerting the appropriate nearby health care professionals when a patient needs care, and
- Providing logistical assistance in responding to medical emergencies.

All of this must be done without incurring undue cost. It is also crucial to not disrupt the life of the patient. The sensors and communication hardware must be lightweight, unobtrusive, and energy efficient. Furthermore, the system must be able to quickly locate the patient in the event of an emergency, without sacrificing the patient's privacy.

A key issue is the design of the alerting system. If it "cries wolf" too often, the response time of medical personnel is bound to decline and the patient will be inclined to stop using it. On the other hand, failure to raise and alert when appropriate could have catastrophic consequences.

A.3. Approach

The approach we propose is to further develop and deploy the Patient Centric Network, Cricket (an indoors location system), INS (a system for finding devices in an environment based on location and other attributes), and an interface to existing BWH systems.

The Patient Centric Network will be expanded to incorporate pulse oximetry and one and two-lead EKG sensors. Algorithms will be developed for analyzing these signals and alerting providers to critical events.

Cricket will be deployed to provide indoor location information. It will be integrated with INS to provide a location database of providers and patients. RFID will be deployed for equipment tracking and integrated with INS.

Software will be provided to interface with provider PDAs and desktop computers to the location database, so that location queries can be resolved.

The monitoring of patients' vital signs is not a new problem: we seek to make it available in a cost-effective, patient-and-provider friendly way in non-traditional environments, especially where there are many mobile patients. The number of false alarms will be reduced by correlating the data from multiple sensors (sensor fusion) to get information about the state of the patient that is both accurate and robust with respect to errors generated by sensors.

Within a medical system, patient data needs to be protected. To this end, we will deploy encrypted network links, use encryption to protect stored data, and restrict access to the data. Further, caregivers, patients, medical devices, and software will need to be authenticated to the system. We plan to use SDSI (Simple Distributed Security Infrastructure [sdsi]) developed at LCS by Prof. Rivest's group, to provide authentication.

As the project develops in scope, more technologies will need to be developed and deployed and strategies for managing larger scale enterprises will become more central. When solving logistics problems, it often is important to search local information first, and then look for less local solutions. In doing this it is often necessary to respect administrative domains. For example, while most ER equipment can be available to any provider within the ER, most ER equipment is not normally available to other departments, except under unusual conditions. We will look to SLAM and Twine [Balazinska], new research initiatives, for approaches to these problems.

A.4. Methods

The following sections outline the base technologies we will use in building the SMART test bed network. Next the basic components are described. We then discuss integrating these components to produce a basic system. Finally we look briefly at expanding the scope of the system beyond the initial deployment and propose strategies for energy management.

A.4.1 Base Technologies

The following subsections describe the base technologies used in building the SMART system. They include the Patient Centric Network, the Cricket System, RFID, and INS. The Patient Centric Network focuses on collecting sensor data from each patient and processing it. The Cricket System provides location information for both patients and healthcare providers. The RFID system provides location information for equipment. INS provides a location database for patient location information, provider location information, and equipment information. These systems are all inherently scalable. We will look to the SLAM and Twine projects to provide further insights into scalability issues.

A.4.1.1 Patient Centric Network

The Patient Centric Network (PCN) is an ongoing research project at MIT. Our goal is to build a flexible, cost-effective patient monitoring system prototype based on commodity hardware, wireless communications and lightweight sensors and actuators.

This project evolved from our SpectrumWare software radio project [Bose 99]. In SpectrumWare we pushed the boundary between software and hardware in the communications domain. We captured a wideband signal, digitized it close to the antenna, moved the samples into memory, and then did all signal processing on a commodity PC. This allowed multiplexing of the same hardware for a variety of applications (ranging from a multi-mode cell phone, to a wireless patch panel, to a television), while simultaneously allowing us to experiment with novel algorithms for transmission and reception.

Here we are using a similar technique in a different domain. A medical device consists of a sensor or actuator connected to an A/D converter and a network interface. Our architecture, Figure 1, is multi-tiered. It employs a collection of gateways and software proxies, running on general-purpose hardware, to bring sensors and actuators onto the network.

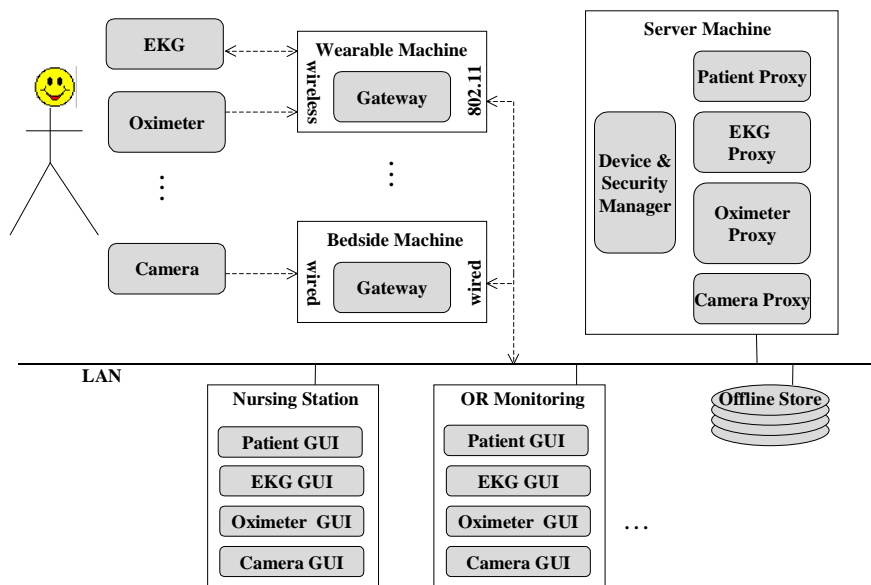


Figure 1: The multi-tiered architecture of our Patient-Centric Network.

Gateways serve as relays to the main network, providing a bridge between the various communication technologies and protocols supported by sensors and the standardized technology used in the rest of the system. Supporting a new connection technology or protocol requires updating only the gateway, not the other network components.

Compared to radios, most medical sensors generate data at a relatively low rate (on the order of Hz or kHz). Furthermore, the sensors communicate with a gateway that is guaranteed to be within close proximity (approximately one meter). Together, these two facts often make it possible to use inexpensive, low power, wireless interfaces between the devices and the gateways. Furthermore, since loss of an occasional sample will not impact the overall functionality or utility of the system, devices need not attempt to retransmit lost packets.

To facilitate our goal of fusing data streams, gateways also time stamp each message before relaying it. This is a critical step. Sensor fusion requires the temporal synchronization of data from multiple sensors. Many sensors will not have a clock, and those clocks that do exist are not likely to be synchronized with each other. In contrast, each gateway has a clock, and these clocks are kept synchronized using a standard algorithm [Mills 92].

Each proxy is an application instance executing on a shared, general-purpose machine. All information processing occurs in the proxies, which can be chained together to perform sophisticated analysis, data fusion, and to create new virtual devices. The proxies perform a full spectrum of information flow operations, from communicating with devices via the gateways, to providing information to user interface applications including sophisticated alerting systems. In general, a proxy receives data from, or sends commands to, one or more upstream components; and serves data to, or accepts commands from, one or more downstream applications.

For every physical device, there is a single proxy that communicates directly with that device via a gateway. These are called device proxies. The device proxy must understand the device's (often proprietary) operating semantics so that it can provide an accessible interface to the device's functions. Our expectation is that device manufacturers will ship a default proxy with the device. Note that all components downstream from a device proxy make no distinction between it and other types of proxies.

A key advantage of using computation proxy chains to process information is the modularity of the construction. New data are readily obtained by replacing individual proxies, rather than larger, monolithic programs. Some of the most interesting proxies combine data from multiple sensors to provide synthetic devices.

User interfaces are completely separated from both devices and proxies. The separation facilitates remote monitoring on devices ranging from large format high-resolution displays to cellular telephones.

Unlike most other networks of devices, e.g., [Gribble 2001] [Heinzelman 2002], our network relies upon a centralized network manager. This software is responsible for managing the network, including detecting the arrival and departure of devices, connecting proxies to devices and to each other, supplying data to remote monitors, and providing security. The network manager is part of the trusted computing base (TCB), and therefore resides in a different administrative domain than other parts of the system. It is responsible for authenticating proxies as well as users.

We rely on a centralized network manager for several reasons. The primary reason is that a central authority most readily detects errors involving collections of components in conflicting states. Since the number of sensors connected to an individual patient will not be huge (probably less than 100), centralized management does not introduce a scaling problem.

A second reason for centralized management is that it conforms to standard hospital operating procedures. The network manager can reside in a separate administrative domain, subject to whatever controls are appropriate.

The network manager's foundation is a static database that describes known device and proxy types. The records include numerous attributes detailing the components' functionality and requirements. These attributes are particularly important when components join the network.

The network manager also tracks dynamic state about device, gateway, and proxy instances currently on the network. The information includes their operating status, along with the connections and dependencies between components.

As the number of patients to be monitored increases, more network managers can be added.

A.4.1.2 Indoor Location System: Cricket

Cricket is an indoor location system for pervasive computing environments. These environments take advantage of emergent network-enabled devices and the promise of ubiquitous network connectivity. A compelling set of applications in pervasive environments is *context-aware*, being able to discover the external context in which they run and adapt accordingly. An important example of context is *location*, such as the position (in some coordinate system) of a device or user, the geographic space in which a device or user is (e.g., the room or portion of a room), and the orientation of a device within some coordinate system. Knowledge of location in the form of coordinate position, spatial resolution, and orientation (a.k.a. "directionality" or "heading") enables a wide variety of pervasive computing applications such as resource discovery, "point-and-use" interfaces and navigation. In the proposed system, the patient monitoring devices and the healthcare providers' devices will update the location database with their Cricket location information. This will enable healthcare providers to find patients and other healthcare providers when a patient is in distress.

While location information in outdoor environments may be obtained via the Global Positioning System or using the cellular infrastructure (with the emerging E-911 services), such capabilities are unavailable in indoor environments or around tall buildings where line-of-sight to GPS [Getting 93] satellites is usually unavailable. However, location-aware applications *inside* buildings, such as offices (and campuses), shopping malls, airports, homes, hospitals, etc. have the potential to fundamentally change the way they interact with their immediate environment. Obtaining location information for applications in an indoor environment in an *unobstrusive* and *private* manner is a challenging task. Indoor environments are harsher than outdoor ones in their treatment of radio signals because of multipath effects and dead spots inside buildings.

User-privacy concerns are an important consideration in the successful deployment of these applications. The administration of the hardware and software infrastructure used for this must be minimal because of the large number of devices and networked services that need this information.

Technology

Cricket uses a combination of Radio Frequency (RF) and ultrasound technologies to provide a location-support service to users and applications. Wall- or ceiling-mounted *beacons*, deployed within a building, publish information on an RF signal in the 418 MHz AM band. With each RF advertisement, a beacon transmits a concurrent ultrasonic pulse. Compact *listeners* attached to mobile or static devices listen for RF signals and, upon receipt of the first few bits, listen for the corresponding ultrasonic pulse. When this pulse arrives, the listeners compute a distance estimate for the corresponding beacon. The listeners run an inference algorithm to correlate the RF and ultrasound signals (the latter are simple pulses with no data encoded on them) emitted from the same beacon. Even in the presence of several competing beacons, our goal is to accurately estimate the linear distances between a listener and all in-range beacons within a small number of seconds.

Cricket uses active beacons and passive listeners, which has three significant benefits. First, it scales well as the number of devices increases; a system with active transmitters attached to devices wouldn't scale particularly well with the density of instrumented devices. Second, its decentralized architecture makes it easy to deploy. This does not mean it is hard to manage; a centralized front-end allows easy management and control. Third, since the listener is passive, the listener controls dissemination of information about its location, thus maintaining privacy

A.4.1.3 RFID

RFID, Radio Frequency Identification [Finkenzeller 99][Sarma 00], is a commercially available system being developed to replace bar codes. As in the bar code system, each item is tagged with a unique identifier. In the bar code system, each identifier is a number. A bar code scanner is used to transfer the number from the item to an application. The application then uses the number to find information about the item from a database. A common example of the use of bar codes is in retail operations, e.g., grocery stores or department stores.

In an RFID system, there is also a tag. This time it is electronically based. An RFID reader is used to transfer the information encoded in the tag to an application. Usually the tag is passive and has no power source (i.e., no battery). The electronics on the tag are activated when the RFID tag is scanned. The scanning process has the reader generating an electromagnetic field and inducing a current in the RFID tag electronics. The current gives power to the RFID tag and allows it to emit its signal. This signal contains the identifying information. As in a bar code system, the application uses this information to select relevant information from a database.

Both systems can be used for inventory control. The chief advantage of an RFID system is that the scanning process is more tolerant of spatial ambiguity, i.e., the reader need not be very close to the tag. RFID scanning can be similar to walking through a metal detector at an airport. This *hands-off* scanning approach of RFID is useful in a hospital environment, where healthcare providers' attention is focussed on patients and not on inventory.

A.4.1.4 Location Database: INS

The Intentional Naming System, INS is designed for naming and discovering a variety of resources in networks of devices and services. INS names are *intentional*; they describe application *intent* in the form of properties and attributes of resources and data, rather than simply network locations of objects (which is the way most traditional network naming systems work today). Names express and satisfy queries about information and their providers. Queries often encode the node doing the query and the results can be exact matches or range matches.

INS resolvers are called Intentional Name Resolvers, or INRs. In any domain, INRs configure themselves into an application-level overlay network based on performance metrics and exchange meta-data about names and the corresponding network locations. The configuration protocols will enable the system to work with no manual intervention, incorporating machinery to bootstrap INRs, spawn and terminate INRs, maintain neighbor relationships, and perform load management across both the local and wide-area Internet.

In addition to achieving the above properties, an important challenge in INS is scaling to a large number ("millions") of names and services. While this kind of scaling will not be relevant to our demonstration systems in the ER, it will be necessary when SMART is transferred to wide area applications. Internet routing relies almost solely on hierarchy to achieve scaling. INS cannot do this, because intentional names aren't necessarily hierarchical. To combat this, INS uses peer-to-peer techniques to scale [Balazinska].

A.4.2 Components

A.4.2.1 Tracking Equipment: RFID

We are planning to use RFID technology to track equipment. Initially each piece of equipment will be tagged with an RFID tag and added to the database of equipment. RFID readers will be added to the ER area. Each reader will have a virtual tag. This allows the reader and its location to be entered into the location database. As a piece of equipment passes an RFID reader, the reader will transmit its identity and the equipment's RFID tag to an RFID reader manager. This reader manager will update the location information associated with that piece of equipment in the location database.

A.4.2.2 Tracking Patients: PDAs

It is important for healthcare providers to be able to find patients. This is especially true if the patient's monitoring information shows that the patient is in need of assistance.

While some people prefer not to be tracked, it is hoped that these patients will tolerate tracking in exchange for better care.

The method of choice in this environment is to provide the patient with a PDA. This PDA will be used to interface to vital signs monitoring sensors and to a Cricket listener. The Cricket listener will listen for messages from Cricket beacons and pass these to the PDA. At the patient's discretion, the PDA will then update the patient's location in the location database. Alternatively, the PDA can be configured to inform the system of the patient's location only when some event, e.g., a medical emergency, arises.

A.4.2.3 Provider PDAs

In emergencies it can be critical for healthcare providers to be able to locate other healthcare providers. We hope that healthcare providers will be comfortable being tracked.

The method will be similar to tracking patients: each healthcare provider will be provided with a PDA. This PDA will provide the healthcare provider access to patient alerts, patient records, and the location of equipment and other healthcare providers. This PDA will interface to a Cricket listener. The Cricket listener will listen to Cricket beacon messages and pass them to the PDA. The PDA will then update the location information associated with the healthcare provider in the location database.

A.4.2.4 Sensors

In support of the SMART system, we will interface PDAs with oximetry sensors and EKG sensors with small numbers of leads. These sensors will provide the ER healthcare providers with improved monitoring of patients who are between consultations or waiting for test results.

For oximetry sensors, we are considering Nonin Xpods. These are lightweight self-contained units with a serial port for transmitting data to the patient's PDA.

We are currently engaged in adding a 2-lead EKG sensor to the Patient Centric Network.

A.4.3 Integrating the Components

Here we review how the previously presented components can be integrated to meet the goals of the SMART system.

A.4.3.1 Patient Monitoring

One of the key goals of this project is to improve monitoring of patients' vital signs while they are in the ER: patients spend a significant amount of time in the ER, not in the presence of a healthcare provider. During this time, critical events can occur and drastic changes can take place in the patient's health status without being visible to ER healthcare providers. SMART will provide cost-effective, patient-acceptable monitoring to prevent these undesirable outcomes.

The process here will be to provide the patient with a PDA. The PDA will interface to various sensors that are monitoring vital signs. The PDA will also interface to the Cricket location system. This will provide the PDA with location information for the patient. The patient's PDA will update the location database with the patient's location. The patient's PDA will timestamp all sensor data and pass it to the sensor proxies on the department computer that will be tracking patient status.

The department computer will have algorithms to determine whether alerts should be forwarded to healthcare providers.

A.4.3.2 Interfacing with the Decision Support System

The decision support system will be critical to making this system work well. Its chief goals are to alert healthcare provider(s) when there is a problem with a patient and to provide healthcare providers with access to information concerning patient status, patient location, equipment location, the location of other healthcare providers and a recommendation about how to proceed.

A.4.3.3 Alerts

Alerting systems play an important role in many clinical settings. They are used to alert medical personnel to conditions that arise in unattended patients. In situations where attending personnel may be subject to "information overload" they are used to direct the attention of medical personnel. We used the prototype PCN discussed above to experiment with a cardiac annunciator panel running over a wireless connection on a handheld device.

Patient Care Network			
Condition	CVP	SBP	HR
Normal HR + RHY + BP		> 100 mmHg	60 < HR < 100 bpm
Tachycardia			> 100 bpm
Bradycardia			< 60 bpm
Junctional rhythm	cannon waves		> 100 bpm
Supraventricular Tachycardia	no change		> 140 bpm
Arrhythmia with low BP		< 100 mmHg	> 100 bpm
Arrhythmia without low BP		> 100 mmHg	> 100 bpm
Mute	SBP :	121 mmHg	HR :
			104 bpm

Figure 2: Annunciator panel used to raise alerts for cardiac conditions. When a condition arises, appropriate part of panel turns red. If the panel is not muted, a verbal alarm is raised as well.

The graphical display, Figure 2, provides alerts for a variety of dangerous cardiac states, and is driven by a collection of proxies that provide signal processing and sensor fusion.

This virtual device uses raw data from three physical devices: central venous pressure (CVP), systolic blood pressure (SBP), and EKG.

To detect backward blood flow, the cannon wave proxy receives the CVP stream. A SBP heart rate proxy connects to the SBP device proxy and performs a heart rate calculation [Akay 94]. Similarly, the EKG rate proxy produces another heart rate calculation from the EKG device proxy. Both rate proxies compute the values using time stamps generated at the gateway. Finally, a robust heart rate proxy fuses the two primary rate values by performing a weighted average on their values. The fused heart rate is robust in that it is derived from multiple physiological signals.

This simple alerting system was built in an *ad hoc* fashion by getting physicians to critique several versions. It is clear to us, however, that this approach will not extend to more complex situations. What is needed is a principled and systematic way to build alerting systems.

We would like to develop a flexible software architecture that links sensors and actuators with patients and caretakers and provides sophisticated mechanisms for detecting problems and raising alerts. A generic, ideal alerting system would detect every problem, issue no false alarms, provide sufficient advanced notification to allow for preventative interventions, and warn medical personnel when an intervention has the potential to preclude potentially useful future treatments.

Of course, a perfect 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.

Given this uncertainty, it seems prudent to design an architecture that carefully separates policy from mechanism. Medical personnel should be able to easily specify whatever alerting policy they desire, and the system should have mechanisms that automatically implement it. Additionally, the system should provide continuous feedback about how well (or poorly) various policies perform, and incorporate technology that allows the system to adapt to changing circumstances. Recently, we have started looking into the design of the cockpit alerting systems deployed in the Traffic Alert and Collision Avoidance System (TCAS) used by commercial aircraft. This environment has a number of things in common with medical environments, and a number of very important differences.

We plan to start our work by designing a formal model of the alerting system for a number of medical environments. It will be based on an approach devised by Jim Kuchar of MIT's Department of Aeronautics and Astronautics [Kuchar 96][Kuchar 01][Song 01]. He 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.

Sensors for patient and airplane monitoring have comparable errors that are amenable to probabilistic modeling. Using those readings to generate an estimate of the current state is quite different. With airplanes, one generates a single state estimate, along with an area around that point that accommodates for errors in the sensor data. In contrast, there are often multiple interpretations of a given physiological reading. Predicting future states is also much more difficult in the medical domain. Knowledge of airplane physics is far more precise than that of human physiology. Therefore, a far greater number of next states exist for patient monitoring. The number of possible interventions is also larger in the medical domain. Whereas a pilot can change an airplane's acceleration, altitude and direction, a doctor can apply a multitude of therapies. Moreover, medical interventions can interact in complex ways. For example, administering one drug might preclude the administration of other relevant drugs. Providing alerts about precluded treatment options is both important and challenging.

A.4.4 Expanding the Scope

The following subsections examine expanding the SMART system within a hospital, to outdoors locations, and its applicability to large scale disaster management.

A.4.4.1 Expanding within the Hospital

Twine [Balazinska] is a resource discovery system that builds on INS, using Chord [Stoica 01] as a distributed hash lookup. It uses the same naming syntax as INS. It offers both discovery and early binding functionality to client applications. Twine implements a new form of intentional name resolution that achieves scalability using a hash-based partitioning of resource descriptions among the Intentional Name Resolvers.

Twine does not require pre-configured hierarchies or special naming syntax. It works with arbitrary attribute sets and achieves balanced resource distribution among participating resolvers. It also handles queries based on orthogonal and hierarchical attributes, with no content or location constraints.

Twine uses a set of resolvers, Twine nodes, that organize themselves into an overlay network [overlaynetdefn] to route resource descriptions to each other for storage, and to collaboratively resolve client queries. Each resolver dynamically specializes in learning about a subset of other Twine nodes, as well as a subset of available resources.

The high-level motivation for Twine's design comes from peer-to-peer document distribution architectures like Freenet[Freenet]. These systems offer interesting scaling possibilities by avoiding central bottlenecks, and by having nodes specialize in subsets of the entire document space. Twine treats resource descriptions the way a system like Freenet might treat an entire document, although the details of how distribution is done are very different. Twine also focuses on resource information distribution for an efficient resolution of queries with *incomplete* resource descriptions. The Twine architecture is layered, building on top of Chord. Chord's simplicity makes it an attractive substrate.

A.4.4.2 Expanding Outdoors

Expanding this system to work outdoors involves interfacing with GPS to get location information and with the cellular network for wide-area communications.

A.4.4.3 Large Scale Disaster Management

When a large-scale disaster occurs, wired communications is often lost. Therefore easily deployable wireless communications systems are quite important. Adding an easily deployed location information system will help rescue workers communicate useful information to each other. Easily deployed patient monitoring systems will help rescue workers in a challenging environment deploy their efforts to best effect. Networks that are self-organizing, sometimes called *ad hoc* networks, are essential in dealing with a chaotic environment.

A.4.5 Energy Management

Currently mobile devices are highly dependent on battery power. Wireless communication is a particularly large consumer of power. Bluetooth [blue99] is more energy efficient than 802.11b. Bluetooth may be deployed between the vital signs monitoring sensors and the gateway. Other technologies may evolve in this area. The component technologies we plan to use have been developed with energy efficient in mind. PCN has been designed to facilitate the use of low power communication at the edges of the network, where battery powered devices are prevalent. Though the current generation of Cricket beacons and listeners consume too much energy, we are in the process of redesigning them in a way that will significantly reduce energy consumption.

Related Work

The Patient Centric Network, a pervasive network of medical sensors, crosses many problem domains and hence shares context with much related work. We divide the work into three broad categories: point solutions for medical sensors, industrial medical or information networks, and other architectures for pervasive sensor networks.

There are an ever increasing number of well publicized point solutions to medical problems that leverage pervasive technology. Articles frequently appear in the technology sections of publications such as *The New York Times* and *The Wall Street Journal* that discuss the use of wireless links and palmtop machines to monitor physiological signals [Gaither 01],[ortivus]. A common application is to monitor EKG on a Palm Pilot or Compaq iPaq [ipaqEKG],[pdaMD]. While each new solution is exciting, they all continue along the traditional path of point solutions in medical technology.

Closer to PCN's goals are information distribution solutions, such as those offered by Tibco [tibco]. Tibco provides an enterprise-level systems integration solution that can meet real-time process demands. While some of their techniques might be applicable

to PCN's data distribution, Tibco's architecture is not suitable for sensor networks. In particular, it does not address issues of interfacing with a large variety of individual devices, nor does it provide for fusing streams of data from sensors.

A system with similar goals to PCN is Frontiers by eko systems, Inc. [ekosystems]. Their vision is an electronic medical record system that eliminates the "enormously labor intensive and potentially error-prone paper process used today in charting surgical patients." The Frontiers system uses devices similar to PCN's gateway to connect devices or other information sources to centralized charting and storage servers. They focus on accessing and retrieving a broad range of data, and do not consider dynamic fusion or processing of data streams.

Another related system is the **Ninja** architecture, developed at Berkeley [Gribble 01],[ninjawww]. That project aims to "develop a software infrastructure to support the next generation of Internet-based applications." Like PCN, Ninja is designed to incorporate simple devices, or *units*, into large, highly available Internet applications. It also has the concept of software proxies that allow composition and customization of applications. It does not deal with the fusion of streams of sensor data. Rather, its goal is to fuse services such as email clients and cell phones in a modular way. In addition, the Ninja system does not have a time-critical mechanism for monitoring the network and responding to changes or errors.

There are other location and tracking systems available: the Bat system [Harter 87], the Active Badge system [Want 92], and RADAR [Bahl 00]. These systems are based on the tracked item sending a signal to sensors that have been deployed on ceilings or walls of the building. The sensor then reports the item's location to a central system. These systems are more difficult to deploy because the sensors need to be wired to a central computer. This also makes them less scalable. This is in contrast to the Cricket location system where the sensors are on the tracked item and the beacons are on the walls or ceilings of the building and are **not** connected to any network.

References

- [Adjie-Winoto 99] Adjie-Winoto, W., Schwartz, E., Balakrishnan, H. and Lilley, J. "The design and implementation of an intentional naming system," In Proc. ACM Symposium on Operating Systems Principles (Kiawah Island, SC, Dec. 1999), pp. 186–201.
- [Akay 94] Akay, Metin, *Biomedical Signal Processing*, Academic Press, 1994.
- [Bahl 00] Bahl, P., Padmanabhan, B. RADAR: An In-Building RF-based User Location and Tracking System. Proc. IEEE INFOCOM (Tel-Aviv, Israel, Mar. 2000)
- [Balazinska] Balazinska, M., and Balakrishnan, H. Twine: Scalable intentional resource discovery for pervasive computing environments. <http://nms.lcs.mit.edu/projects/twine>.
- [Bose 99] V. Bose and M. Ismert and M. Welborn and J. Guttag, "Virtual Radios," IEEE Journal on Selected Areas in Communications, vol. 17, no. 9, April 1999
- [blue99] Specification of Bluetooth System, <http://www.bluetooth.com>, December, 1999.
- [ekosystems] eko systems, Inc. Frontier, electronic information management. <http://www.ekosystems.com>

- [Finkenzeller 99] Finkenzeller, K. *RFID Handbook—radio-frequency identification fundamentals and applications*. John Wiley & Sons, 1999.
- [Gaither 01] Gaither, Chris. “Bluetooth defies obituaries.” In *The New York Times*, December, 2001.
- [Gribble 01] Gribble, Steven, “The Ninja Architecture for Robust {Internet}-scale Systems and Services,” *Computer Networks*, April 2001.
- [Harter 97] Harter, A., Hopper, A., “A New Location Technique for the Active Office.” *IEEE Personal Communications* 4, 5 (October 1997), 43-47.
- [Harfst 02] Harfst, Gregory and Guttag, John, “A Patient-Centric Network: An Architecture for Pervasive Medical Devices,” submitted for publication.
- [Heinzelman 02] W. Heinzelman and A. Chandrakasan and H. Balakrishnan, “An Application-Specific Protocol Architecture for Wireless Microsensor Network,” *IEEE Transactions on Wireless Communications*, to appear 2002.
- [ipaqEKG] MicroMedical. MicroMedical PocketView ECG. http://www.micromed.com.au/08_products/20_pocketview/
- [Kuchar 96] Kuchar, James K., “Methodology for Alerting-System Performance Evaluation,” *Journal of Guidance, Controls, and Dynamics*, Volume 19, Number 2, Mar-Apr 1996.
- [Kuchar 01] Kuchar, James K., “Managing Uncertainty in Decision-Aiding and Alerting System Design,” *6th CNS/ATM Conference*, March 2001.
- [Mills 92] Mill, David, “Network Time Protocol (Version 3) Specification, Implementation and Analysis,” IETF RFC 1305, March 1992.
- [ninjawww] Ninja group at Berkeley. Ninja homepage. <http://ninja.cs.berkeley.edu>.
- [overlaynetdefn] A network is a collection of interconnected nodes. An overlay network is based on a network and uses a subset of the base network’s nodes and a subset of the base network’s connections. The method of choosing which nodes and which connections are in the overlay network depends on the purpose of the overlay network.
- [Oxygen] Project Oxygen, a Consortium at MIT <http://oxygen.lcs.mit.edu>
- [pdaMD] pdaMD.com <http://www.pdamd.com/vertical/home.xml>
- [Priyantha 00] Priyantha, N., Chakraborty, A., and Balakrishnan, H. “The Cricket Location-Support System.” In Proc. 6th ACM MOBICOM Conf. (Boston, MA, Aug. 2000), pp. 32–43.
- [ortivus] Ortivus, MobiMed 200. <http://www.ortivus.com/>
- [Sarma 00] Sarma, S., Brock, D., and Ashton, K. “The networked physical world: Proposals for engineering the next generation of computing, commerce and automatic identification.” Tech. Rep. MIT AUTO-ID-WH-001, MIT Auto-ID Center, Dec., 2000.
- [sdsi] <http://theory.lcs.mit.edu/~cis/sdsi.html>.
- [Song 01] Song, Lixia and Kuchar, James K., “Describing, Predicting, and Mitigating Dissonance Between Alerting Systems,” *International Workshop on Human Error, Safety, and System Development*, June 2001.

[Stoica 01] Stoica, I., Morris, R., Karger, D., Kaashoek, M. F., and Balakrishnan, H. "Chord: A scalable peer-to-peer lookup service for internet applications." In Proc. ACM SIGCOMM (San Diego, Aug. 2001).

[tibco] Tibco. <http://www.tibco.com>

[Want 92] Want, R., Hopper, A., Falcao, V., and Biggins, J. "The Active Badge Location System. ACM Transactions on Information Systems 10, 1(January 1992), 91-102.