

SYSTEM ARCHITECTURE OF A WIRELESS BODY AREA SENSOR NETWORK FOR UBIQUITOUS HEALTH MONITORING

CHRIS OTTO, ALEKSANDAR MILENKOVIĆ, COREY SANDERS, EMIL JOVANOVIĆ

University of Alabama in Huntsville

chrisaotto@yahoo.com, {milenka | jovanov}@ece.uah.edu, sanderscorey@yahoo.com

Received October 1, 2005

Revised January 10, 2006

Recent technological advances in sensors, low-power microelectronics and miniaturization, and wireless networking enabled the design and proliferation of wireless sensor networks capable of autonomously monitoring and controlling environments. One of the most promising applications of sensor networks is for human health monitoring. A number of tiny wireless sensors, strategically placed on the human body, create a wireless body area network that can monitor various vital signs, providing real-time feedback to the user and medical personnel. The wireless body area networks promise to revolutionize health monitoring. However, designers of such systems face a number of challenging tasks, as they need to address often quite conflicting requirements for size, operating time, precision, and reliability.

In this paper we present hardware and software architecture of a working wireless sensor network system for ambulatory health status monitoring. The system consists of multiple sensor nodes that monitor body motion and heart activity, a network coordinator, and a personal server running on a personal digital assistant or a personal computer.

Key words: Wireless sensors, body area networks, health monitoring, wearable computing.

Communicated by: A Ganz & R Istepanian

1 Introduction

Recent technological advances in wireless networking, microelectronics integration and miniaturization, sensors, and the Internet allow us to fundamentally modernize and change the way health care services are deployed and delivered. Focus on prevention and early detection of disease or optimal maintenance of chronic conditions promise to augment existing health care systems that are mostly structured and optimized for reacting to crisis and managing illness rather than wellness [6].

The anticipated change and emerging new services are well-timed to help cope with the imminent crisis in the health care systems caused by current economic, social, and demographic trends. The overall health care expenditures in the United States reached \$1.8 trillion in 2004, though almost 45 million Americans do not have health insurance [15]. On the other hand, many companies have already been plagued by high-rising costs of healthcare liabilities. With current trends in healthcare costs, it is projected that the total health care expenditures will reach almost 20% of the Gross Domestic Product (GDP) in less than 10 years from now, threatening the wellbeing of the entire economy. The demographic trends are indicating two significant phenomena: an aging population due to increased life expectancy and Baby Boomers demographic peak. Life expectancy has significantly increased from 49 years in 1901 to 77.6 years in 2003. According to the U.S. Bureau of the Census,

the number of elderly over age 65 is expected to double from 35 million to nearly 70 million by 2025 when the youngest Baby Boomers retire [21]. This trend is global, so the worldwide population over age 65 is expected to more than double from 357 million in 1990 to 761 million in 2025. These statistics underscore the need for more scalable and more affordable health care solutions.

Wearable systems for continuous health monitoring are a key technology in helping the transition to more proactive and affordable healthcare. They allow an individual to closely monitor changes in her or his vital signs and provide feedback to help maintain an optimal health status. If integrated into a telemedical system, these systems can even alert medical personnel when life-threatening changes occur. In addition, the wearable systems can be used for health monitoring of patients in ambulatory settings [7]. For example, they can be used as a part of a diagnostic procedure, optimal maintenance of a chronic condition, a supervised recovery from an acute event or surgical procedure, to monitor adherence to treatment guidelines (e.g., regular cardiovascular exercise), or to monitor effects of drug therapy.

During the last few years there has been a significant increase in the number and variety of wearable health monitoring devices, ranging from simple pulse monitors, activity monitors, and portable Holter monitors, to sophisticated and expensive implantable sensors. However, wider acceptance of the existing systems is still limited by the following important restrictions. Traditionally, personal medical monitoring systems, such as Holter monitors, have been used only to collect data. Data processing and analysis are performed offline, making such devices impractical for continual monitoring and early detection of medical disorders. Systems with multiple sensors for physical rehabilitation often feature unwieldy wires between the sensors and the monitoring system. These wires may limit the patient's activity and level of comfort and thus negatively influence the measured results [12]. In addition, individual sensors often operate as stand-alone systems and usually do not offer flexibility and integration with third-party devices. Finally, the existing systems are rarely made affordable.

One of the most promising approaches in building wearable health monitoring systems utilizes emerging wireless body area networks (WBANs) [8]. A WBAN consists of multiple sensor nodes, each capable of sampling, processing, and communicating one or more vital signs (heart rate, blood pressure, oxygen saturation, activity) or environmental parameters (location, temperature, humidity, light). Typically, these sensors are placed strategically on the human body as tiny patches or hidden in users' clothes allowing ubiquitous health monitoring in their native environment for extended periods of time.

A number of recent research efforts focus on wearable systems for health monitoring. Researchers at the MIT Media Lab have developed MIThril, a wearable computing platform compatible with both custom and off-the-shelf sensors. The MIThril includes ECG, skin temperature, and galvanic skin response (GSR) sensors. In addition, they demonstrated step and gait analysis using 3-axis accelerometers, rate gyros, and pressure sensors [18]. MIThril is being used to research human behaviour recognition and to create context-aware computing interfaces [5]. CodeBlue, a Harvard University research project, is also focused on developing wireless sensor networks for medical applications. They have developed wireless pulse oximeter sensors, wireless ECG sensors, and tri-axial accelerometer motion sensors. Using these sensors, they have demonstrated the formation of ad-hoc networks. The sensors, when outfitted on patients in hospitals or disaster environments, use the ad-hoc networks to transmit vital signs to healthcare givers, facilitating automatic vital sign collection and real-time triage [19,9].

In this paper we describe a general WBAN architecture and how it can be integrated into a broader telemedical system. To explore feasibility of the proposed system and address open issues we have designed a prototype WBAN that consists of a personal server, implemented on a personal digital assistant (PDA) or personal computer (PC), and physiological sensors, implemented using off-the-shelf sensor platforms and custom-built sensor boards. The WBAN includes several motion sensors that monitor the user's overall activity and an ECG sensor for monitoring heart activity. We describe the hardware and software organization of the WBAN prototype.

The rest of the paper is organized as follows. Section 2 outlines the general WBAN architecture, defines the role of each component, and describes its integration into a broader telemedical system. Section 3 presents a case study, walking through a typical system deployment and its use. Section 4 describes the hardware architecture. Section 5 details the software architecture of the personal server and sensor nodes, and introduces energy-efficient WBAN communication protocol. Section 6 concludes the paper and discusses possible future research directions.

2 System Architecture

The proposed wireless body area sensor network for health monitoring integrated into a broader multi-tier telemedicine system is illustrated in Figure 1. The telemedical system spans a network comprised of individual health monitoring systems that connect through the Internet to a medical server tier that resides at the top of this hierarchy. The top tier, centered on a medical server, is optimized to service hundreds or thousands of individual users, and encompasses a complex network of interconnected services, medical personnel, and healthcare professionals. Each user wears a number of sensor nodes that are strategically placed on her body. The primary functions of these sensor nodes are to unobtrusively sample vital signs and transfer the relevant data to a personal server through wireless personal network implemented using ZigBee (802.15.4) or Bluetooth (802.15.1). The personal server, implemented on a personal digital assistant (PDA), cell phone, or home personal computer, sets up and controls the WBAN, provides graphical or audio interface to the user, and transfers the information about health status to the medical server through the Internet or mobile telephone networks (e.g., GPRS, 3G).

The medical server keeps electronic medical records of registered users and provides various services to the users, medical personnel, and informal caregivers. It is the responsibility of the medical server to authenticate users, accept health monitoring session uploads, format and insert this session data into corresponding medical records, analyze the data patterns, recognize serious health anomalies in order to contact emergency care givers, and forward new instructions to the users, such as physician prescribed exercises. The patient's physician can access the data from his/her office via the Internet and examine it to ensure the patient is within expected health metrics (heart rate, blood pressure, activity), ensure that the patient is responding to a given treatment or that a patient has been performing the given exercises. A server agent may inspect the uploaded data and create an alert in the case of a potential medical condition. The large amount of data collected through these services can also be utilized for knowledge discovery through data mining. Integration of the collected data into research databases and quantitative analysis of conditions and patterns could prove invaluable to researchers trying to link symptoms and diagnoses with historical changes in health status, physiological data, or other parameters (e.g., gender, age, weight). In a similar way this infrastructure could significantly contribute to monitoring and studying of drug therapy effects.

The second tier is the personal server that interfaces WBAN sensor nodes, provides the graphical user interface, and communicates with services at the top tier. The personal server is typically

implemented on a PDA or a cell phone, but alternatively can run on a home personal computer. This is particularly convenient for in-home monitoring of elderly patients. The personal server interfaces the WBAN nodes through a network coordinator (nc) that implements ZigBee or Bluetooth connectivity. To communicate to the medical server, the personal server employs mobile telephone networks (2G, GPRS, 3G) or WLANs to reach an Internet access point.

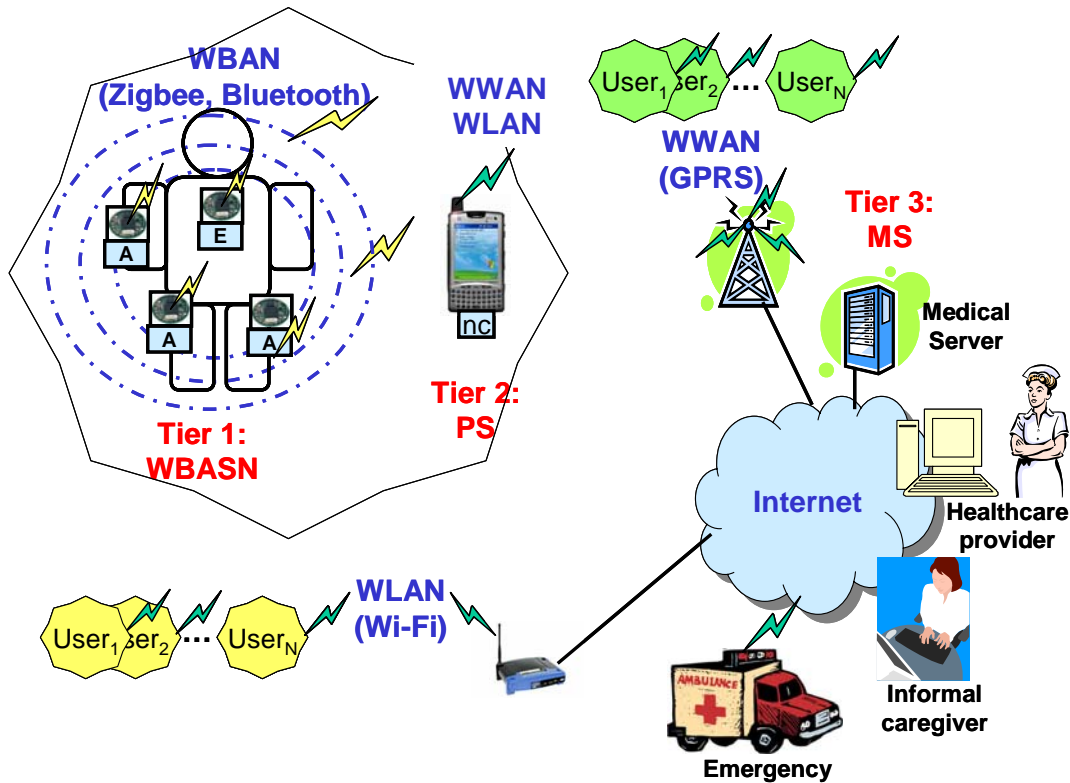


Figure 1 Health Monitoring System Network Architecture

The interface to the WBAN includes the network configuration and management. The network configuration encompasses the following tasks: sensor node registration (type and number of sensors), initialization (e.g., specify sampling frequency and mode of operation), customization (e.g., run user-specific calibration or user-specific signal processing procedure upload), and setup of a secure communication (key exchange). Once the WBAN network is configured, the personal server manages the network, taking care of channel sharing, time synchronization, data retrieval and processing, and fusion of the data. Based on synergy of information from multiple medical sensors the PS application should determine the user's state and his or her health status and provide feedback through a user-friendly and intuitive graphical or audio user interface.

The personal server holds patient authentication information and is configured with the medical server IP address in order to interface the medical services. If the communication channel to the medical server is available, the PS establishes a secure communication to the medical server and sends reports that can be integrated into the user's medical record. However, if a link between the PS and the medical server is not available, the PS should be able to store the data locally and initiate data

uploads when a link becomes available. This organization allows full mobility of users with secure and near real time health information uploads.

A pivotal part of the telemedical system is tier 1 – wireless body area sensor network. It comprises a number of intelligent nodes, each capable of sensing, sampling, processing, and communicating of physiological signals. For example, an ECG sensor can be used for monitoring heart activity, an EMG sensor for monitoring muscle activity, an EEG sensor for monitoring brain electrical activity, a blood pressure sensor for monitoring blood pressure, a tilt sensor for monitoring trunk position, and a breathing sensor for monitoring respiration, while the motion sensors can be used to discriminate the user's status and estimate her or his level of activity.

Each sensor node receives initialization commands and responds to queries from the personal server. WBAN nodes must satisfy requirements for minimal weight, miniature form-factor, low-power consumption to permit prolonged ubiquitous monitoring, seamless integration into a WBAN, standards based interface protocols, and patient-specific calibration, tuning, and customization. The wireless network nodes can be implemented as tiny patches or incorporated into clothes or shoes. The network nodes continuously collect and process raw information, store them locally, and send processed event notifications to the personal server. The type and nature of a healthcare application will determine the frequency of relevant events (sampling, processing, storing, and communicating). Ideally, sensors periodically transmit their status and events, therefore significantly reducing power consumption and extending battery life. When local analysis of data is inconclusive or indicates an emergency situation, the upper level in the hierarchy can issue a request to transfer raw signals to the next tier of the network.

Patient privacy, an outstanding issue and a requirement by law, must be addressed at all tiers in the healthcare system. Data transfers between a user's personal server and the medical server require encryption of all sensitive information related to the personal health [22]. Before possible integration of the data into research databases, all records must be stripped of all information that can tie it to a particular user. The limited range of wireless communications partially addresses security within WBAN; however, the messages can be encrypted using either software or hardware techniques. Some wireless sensor platforms have already provided a low power hardware encryption solution for ZigBee communications [14].

3 Case Study

In this section we present a hypothetical case study to illustrate the usefulness of our proposed system. The patient presented is fictitious, but representative of common issues a recovering heart attack patient would face. We discuss the issues and describe how our system can be used to both address the problem and provide advantages over typical present day solutions.

Juan Lopez is recovering from a heart attack. After the release from the hospital he attended supervised physical rehabilitation for several weeks. His physicians prescribed an exercise regime at home. During the physical rehabilitation it was easy to monitor Juan and verify he completed his exercises. Sadly, when left to his own self-discipline, he does not rigorously follow the exercise as prescribed. He exercises, but is not honest to himself (or his physician) as to the intensity and duration of the exercise. As a result, Juan's recovery is slower than expected which raises concerns about his health prognosis, and his physician has no quantitative way to verify Juan's adherence to the program.

Our health monitoring system offers a solution for Juan. Equipped with a WBAN, tiny sensors provide constant observation of vital statistics, estimate induced energy expenditure, and assist Juan's

exercise. Tiny electronic inertial sensors measure movement while electrodes on the chest can measure Juan's heart activity. The time, duration, and level of intensity of the exercise can be determined by calculating an estimate of energy expenditure from the motion sensors. Through the Internet, his physician can collect and review data, verify Juan is exercising regularly, issue new prescribed exercises, adjust data threshold values, and schedule office visits. Juan's physician need not rely on Juan's testament, but can quantify his level and duration of exercise. In addition, Juan's parameters of heart rate variability provide a direct measure of his physiological response to the exercise serving as an in-home stress test. Substituting these remote stress tests and data collection for in-office tests, Juan's physician reduces the number of office visits. This cuts healthcare costs and makes better use of the physician's time. In urgent cases, however, the personal server can directly contact Emergency Medical Services (EMS) if the user subscribes to this service. Figure 2 illustrates one possible data flow.

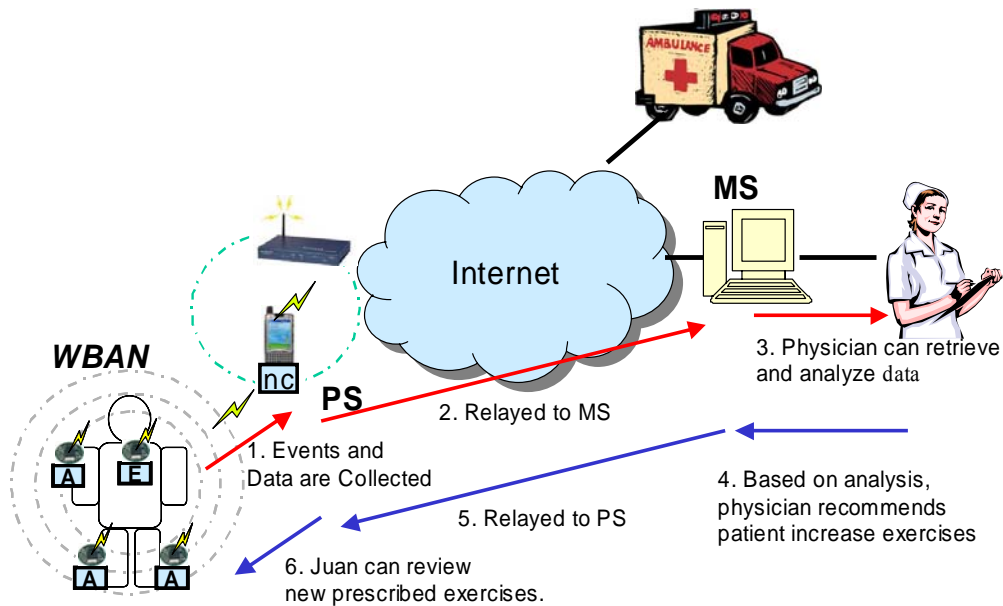


Figure 2 Data flow in the proposed healthcare monitoring system

4 Hardware Architecture

In the spirit of the system architecture presented in Section 2, we have developed a prototype healthcare monitoring system. Figure 3 shows a photograph of the prototype components. The fully operational prototype system includes two activity sensors (ActiS), an integrated ECG and tilt sensor (eActiS), and a personal server. Each sensor node includes a custom application specific board and uses the Tmote sky platform [16] for processing and ZigBee wireless communication. The personal server runs either on a laptop computer or a WLAN/WWAN-enabled handheld PocketPC. The network coordinator with wireless ZigBee interface is implemented on another Tmote sky that connects to the personal server through a USB interface. For an alternative setting we have developed a custom network coordinator that features the ZigBee wireless interface, an ARM processor, and a compact flash interface to the personal server (Figure 3).

The Tmote sky from Moteiv acts as the primary embedded platform for all sensors in our system. Each Tmote sky board utilizes Texas Instrument's MSP430F1611 microcontroller and Chipcon's

CC2420 radio interface. The microcontroller is based around a 16-bit RISC core integrated with 10 KB of RAM and 48 KB of flash memory, analog and digital peripherals, and a flexible clock subsystem. It supports several low-power operating modes and consumes as low as 1 μ A in standby mode; it also has very fast wake up time of 6 μ s. The CC2420 wireless transceiver is IEEE 802.15.4 compliant and has programmable output power, maximum data rate of 250 Kbps, and hardware support for error correction and 128-bit encryption. The CC2420 is controlled by the MSP430 microcontroller through the Serial Peripheral Interface (SPI) port and a series of digital I/O lines with interrupt capabilities. The Tmote sky platform features a 10-pin expansion connector with one Universal Asynchronous Receiver Transmitter (UART) and one I²C interface, two general-purpose I/O lines, and three analog input lines.

The activity sensor, ActiS, consists of the Tmote sky platform and an Intelligent Activity Sensor (IAS), implemented as a daughter card. The IAS monitors motion using two dual-axis accelerometers arranged to provide three orthogonal motion axes (X, Y, Z). The IAS utilizes an on-board MSP430F1232 microcontroller for pre-processing and filtering of sampled data. The IAS connects to the Tmote sky platform via the extension header and sends digital sensor data using a simple serial communication protocol.




Figure 3 Prototype WBAN. From left to right:
the Personal Server with Network Coordinator, ECG sensor with electrodes, and a motion sensor.

Similarly, the integrated ECG and tilt sensor (eActiS) consists of the Tmote sky platform and an intelligent signal processing module (ISPM). The ISPM is similar to the ActiS IAS board, but includes

a single-channel bio-amplifier for three-lead ECG/EMG. Electrodes are connected and placed on the chest for monitoring heart activity. When the sensor is worn on the chest, it also serves as an upper body tilt sensor.

5 Software Architecture

In this section we describe the software architecture of the prototype WBAN system, illustrated in Figure 4. It encompasses software modules running on the IAS/ISMP, the Tmote sky platform, the network coordinator, and the personal server. Our focus has been on developing solutions for real-time on-sensor processing, WBAN communications, time synchronization [16], maximizing battery life [13], managing data and events, and an easy to use user interface. These issues relate to the lower tiers of the network, and as such we describe our prototype software for the WBAN.

5.1 Sensor Node Software

The sensor node software samples and collects physiological data, analyzes the signals in real-time, and transmits the results wirelessly to the personal server. In our prototype this software runs on the Tmote sky platform and custom application specific daughter cards. We have developed software for two types of sensors. An Activity Sensor (ActiS) samples three-axis accelerometers to determine orientation, type of activity (walking, sitting, etc.), estimates activity induced energy expenditure (AEE) based on an algorithm proposed by Bouten, et. al. [1], and performs step detection in real-time. An ECG and tilt sensor (eActiS) monitors heart activity and samples a two-axis accelerometer for orientation (upper body tilt). Sensor node and network coordinator software is implemented in the TinyOS environment.

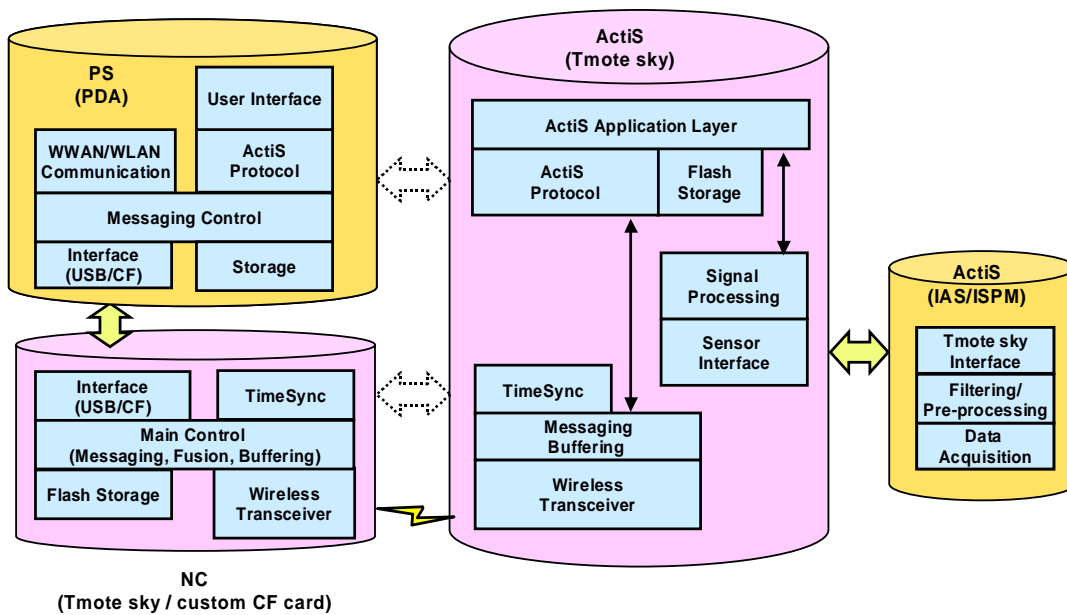


Figure 4 Block Diagram of Software Components in a WBAN.

TinyOS components

TinyOS is a lightweight open source operating system for wireless embedded sensors. It is designed to use minimal resources and its configuration is defined at compile time by combining components from the TinyOS library and custom-developed components. Well-defined interfaces are used to connect and define the data flow between components. A TinyOS application is implemented as a set of component modules written in nesC [20]. The nesC language extends the C language with new support for task synchronization and task management. This approach results in a natural modular design, minimal use of resources, and short development cycles.

TinyOS fully supports the Tmote sky platform and includes library components for the Chipcon CC2420 radio drivers and other on-chip peripherals. Radio configuration, MAC layer communications, and generic packet handling are also natively supported.

Figure 5 depicts the ActiS software architecture using the TinyOS component model. The components in yellow are those reused from the TinyOS library. *GenericComm* provides generic packet handling and basic *SendMsg*, *ReceiveMsg* interfaces using TinyOS messages. A TinyOS message is a generic message structure with a reserved payload for application data [20]. *GenericComm* would also interface to low-level platform specific TinyOS hardware drivers. The ActiS application layer consists of custom components shown in blue.

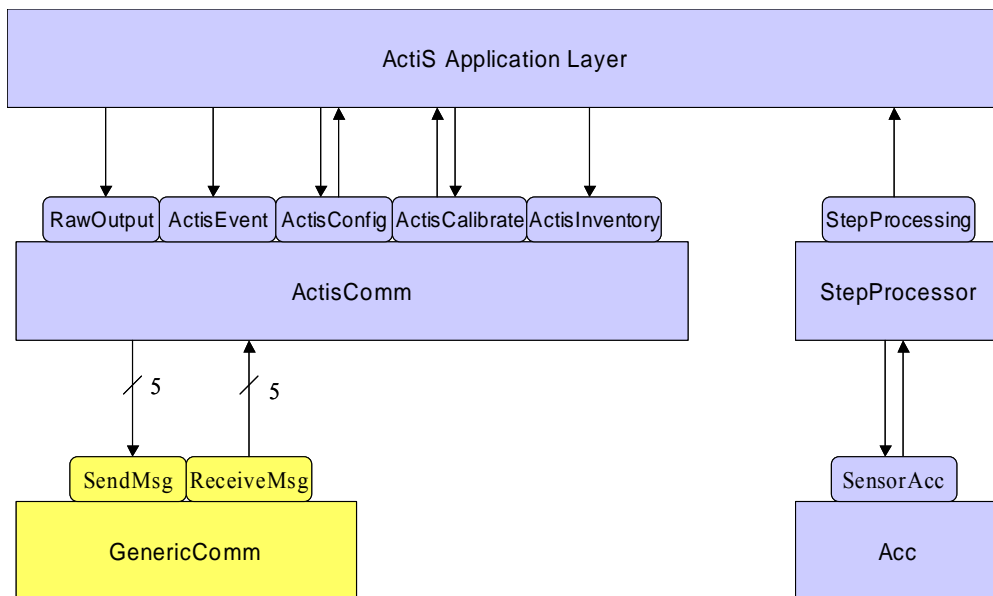


Figure 5 Simplified sensor node interface connection (ActiS).

Communication Protocol and Time Synchronization

Our communication protocol was designed to minimize resources and uphold the spirit of the ZigBee star network topology [24]. Figure 6 shows our communication super frame. All communications are between a sensor node and the network coordinator. Each communication super frame is divided into 50ms timeslots used for message transmissions. Each sensor uses its corresponding timeslot to transmit sensor data, command acknowledgements, and event messages. The first timeslot, however,

belongs to the network coordinator and is used for transmitting configuration commands from the personal server. The network coordinator also transmits periodic beacon messages used to synchronize the start of super frames. This organization also serves as practical collision avoidance, making more efficient use of the available bandwidth when compared to using only CC2420 Collision Sense Multiple Access (CSMA) scheme.

Time synchronization is crucial in providing a means for network communication protocol, as well as for event correlation. In the WBAN, sensor nodes are distributed about the user's body and are only wirelessly connected; the sensors operate for extended periods, sampling and analyzing physiological data. It becomes necessary to correlate detected events between sensors within the sampling interval. A time stamp mechanism can be employed; however, without a global time reference the timestamp has no meaning outside of the scope of a sensor node. Synchronizing the session start times and utilizing a local time reference is not sufficient for the problem at hand. Even if two sensors could precisely agree on the start of a health monitoring session, a running local time would only work for short session durations. Each sensor in the WBAN has a local clock source with an associated skew. The skew is a measure of the difference in frequency between the local clock source and an ideal clock source. As a result of skew, any elapsed time based on a local clock, over time, will differ between any two sensors. This error is cumulative. Consider two 32 KHz crystals differing by a typical 50 ppm (parts per million). Over the course of a few hours, the sensor clocks can differ by more than several hundred milliseconds. For an accurate healthcare monitoring system and correlating step and gait analysis between two sensors this is unacceptable, while the optimization of communication channel sharing would be impossible.




Figure 6 Communication Super Frame and activity of Sensor #2

We address the issue by employing a modified version of the Flooding Time Synchronization Protocol (FTSP) developed at Vanderbilt University [11]. FTSP generates time synchronization by dynamically electing a master node. The master node transmits periodic beacons containing global time stamps. FTSP features MAC layer time stamping for increased precision and skew compensation with linear regression to account for clock drift. We modified FTSP for use in our prototype WBAN [4]. Our modified version exploits the WBAN's star network topology. Beacon messages, used to delineate the super frame, also serve to distribute the global timestamps. Sensor nodes in the WBAN use beacons as a timing reference. For testing of the time synchronization protocol, we developed a test bed where the network coordinator and WBAN sensor nodes are all connected to a common wired

signal. Sensors measured the signal changes in jiffy ticks, where one jiffy is 30.5 μ s, determined by the clock frequency of the on-board crystal (32.768 KHz). In most cases the node's error was within ± 1 jiffy and the average error was approximately $0.1 * T_{\text{jiffy}}$ or 3 μ s.

Power Management

Long-life, persistent sensor nodes require efficient power management. Ease of use and the perceived unobtrusiveness is affected by sensor weight, interval between battery changes, and the level of user interaction required. Because battery life is proportional to battery size (weight) it is our challenge as designers to minimize sensor power consumption and thus maximize battery life for a selected battery size. In designing our prototype we have held low power consumption as a major design goal – both in processor and technology selection as well as software organization. We have selected the MSP430 microprocessor family for their excellent MIPS/mW ratio and 802.15.4 in part because of its unique fit for low power, low data rate applications. Beyond this, it is possible to extend each node's lifetime by clever network organization and making trade-offs between communication and on-sensor computation.




Figure 7 ActiS current consumption during 5 seconds of operation. Red arrow marks one super frame cycle.

Power efficiency can be improved by noting that our wireless sensors consume approximately 22mA when the radio is active and about 1mA when the radio is disabled. Therefore, 95% of the sensors power consumption can be attributed to wireless communications, but it is apparent from our communication protocol that any given sensor node is actively involved in communications for only

two of the available timeslots (see Figure 6). For a one second super frame, the radio is utilized only 10% of the time. By disabling the radio during inactive timeslots, an average current consumption of just 3.1mA can be realized, providing 7 times longer battery life.

Our highly accurate time synchronization makes this possible. Between active time slots the radio is disabled. Immediately before the beacon arrival time, a scheduled timer wakes the processor in order to enable the radio. The beacon is received and then the radio is returned to a disabled state. When a sensor's scheduled timeslot arrives and queued events or data are waiting for transmission, the radio can once again be enabled for transmission. Figure 7 shows the power profiles recorded for a motion sensor using an environment for real time power monitoring [13].

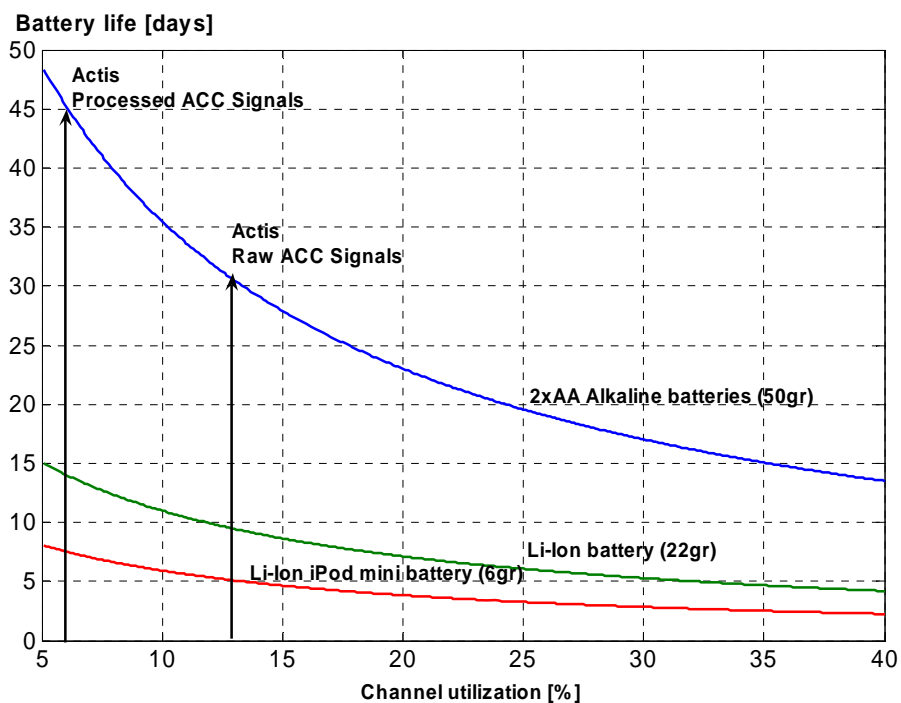


Figure 8 Battery life as a function of channel utilization

It should be noted that the super frame period and timeslot length directly affect power consumption and battery life. Extending the period between beacons allows the radio to stay in low power mode longer. This is not without cost, however. Extending the super frame period also increases the maximum latency between sensor communications and increases the event reporting latency. Every application will have a practical limit of the tolerable event detection latency. Detecting a heart attack may be of little use if the sensor delays notification for up to one hour. For an application specific maximum latency, system designers can make trade offs between battery capacity (weight) and battery life. For our prototype WBAN, a one second super frame period and 50ms timeslots were selected, defining a one second maximum event latency and support for 20 nodes in the network (19 sensors and one network coordinator). Our chosen timeslot and super frame length will effectively create a 5% utilization for individual sensors, as represented in Figure 7. Smart

implementation, however, allows the radio to be disabled as soon as communications are complete – even if the entire timeslot has not expired. In this sense, the timeslot length defines the peak channel utilization for a sensor in the network. Figure 8 shows how sensor battery life can now be described as a function of the type of battery and effective channel utilization.

Power consumption is also conserved by employing on sensor, real time signal processing. Consider a modest WBAN consisting of two tri-axis accelerometers operating at a 200Hz sampling frequency (16-bit samples) and an ECG sensor operating at 1000Hz (16-bit samples). The aggregate data rate, or minimum bandwidth the system must support, is:

$$BW = 2 \text{ sensors} \times 200\text{Hz} \times 3 \text{ axes} \times 16 + 1000\text{Hz} \times 16 = 35.2\text{kbps}$$

Application bandwidth (radio utilization) directly influences power consumption. Eliminating raw data transmission and instead transmitting pertinent events can save significant bandwidth, but necessitates an event management scheme.

Event Management

In the WBAN prototype a sensor node generates an event when a characteristic feature has been recognized. An event is described by event type, timestamp, and context-relevant data. An ECG sensor processes raw ECG signal and detects R-peak event in QRS complex. A motion sensor can detect step events, change in user orientation (sitting to standing), or AEE (activity induced energy estimates) events providing a measure of exercise intensity.

Sensor node software is responsible for detecting the events in real-time. The current implementation of the prototype supports R-peak detection using a modified version of the algorithm presented by Pan, et. al. [17], and user activity (AEE) events based on an algorithm proposed by Bouten, et. al. [1]. ST segment processing can be implemented using an approach as presented by Maglaveras [10] and Wheelock [23]. Step detection algorithms are still in development.

For a health monitoring WBAN such as ours, it is imperative that event data can be correlated among the sensors in the distributed sensor network. Time synchronization addresses this issue. Each event message contains a 32-bit timestamp in “jiffy” ticks indicating the precise event occurrence relative to the start of a health monitoring session.

Figure 9 shows how events are processed in a WBAN. Events occur asynchronously relative to a sensor’s designated timeslot and must be queued for transmission up to one super frame period (one second for our WBAN). At the precise time of event occurrence, an event message with global timestamp is constructed and the message is queued for transmission. When the sensor’s scheduled timeslot arrives, all pending messages are transmitted. Currently, we have implemented a redundant data transmission scheme where each event message describes the current event and a previous event, making the system somewhat resilient to packet loss errors. For particularly critical events, however, an explicit acknowledgement is desired to ensure arrival at Personal Server. Each packet also includes a unique frame sequence making it possible to retransmit event messages in the next timeslot if necessary. This buffering mechanism is limited by available memory; for the Tmote sky platform, up to 160 messages can be buffered. Although in practice, no more than 25 message buffers are needed for the events of interest (heartbeat, steps, etc.).

Both the ActiS and eActis platforms require about 25 KB of flash memory for TinyOS and program space and about 1.1 KB RAM for data and an additional 1.4 KB for message buffering (25 buffers with 56 bytes per message). It should be noted that changing the super frame period (discussed in section 5) and increasing latency will also define the required number of message buffers for an

application. This yields 52% flash utilization and 25% RAM utilization for a sensor node implemented on a Tmote sky platform.




Figure 9 Events are generated and queued for transmission during the sensor’s scheduled timeslot.

5.2 Personal Server Software

The Personal Server provides the user interface, controls the WBAN, fuses data and events, and creates unique session archive files. The software is implemented in Visual Basic using Visual Studio .NET 2003. It runs on either a Windows CE Pocket PC or a Windows PC. User archive files are created in real-time using a custom binary file format, and are then converted to a Microsoft Access database for off-line analysis. Based on our communication protocol and size of the super frame, we can support as many as nineteen sensor nodes, although in practice our typical WBAN research involves only three: two motion sensors and one heart rate sensor.

Figure 10 shows the message flow during a typical WBAN health monitoring session. The Personal Server begins a health monitoring session by wirelessly configuring sensor parameters, such as sampling rate, selection of the type of physiological signal of interest, and specifying events of interest. For example, our ActiS motion sensors are capable of step detection if placed on the ankles, upper body tilt if placed on the chest, and contribute differently to energy estimation depending on location. Sensors in turn, transmit pertinent event messages to the personal server. The personal server must aggregate the multiple data streams, create session files and archive the information in the patient database. Real-time feedback is provided through the user interface. The user can monitor his / her vital signs and be notified of any detected warnings or alerts.

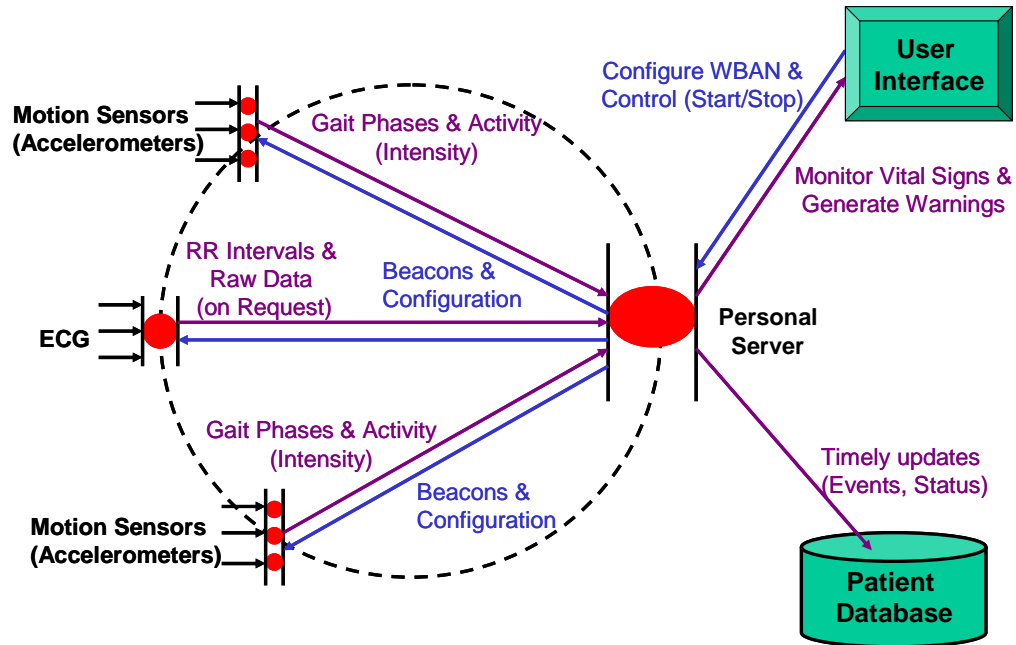


Figure 10 Data flow in prototype WBAN

Graphical User Interface

We designed the user interface to both address the needs of the research prototype WBAN and support a deployed WBAN system. The user interface must provide seamless control of the WBAN, implementing all the necessary control over the WBAN. In designing the user interface we identified five design goals that the PS must support:

- Node Identification
- Configuration of sensors
- Calibration of sensors
- Graphical presentation of Events and Alerts
- Visual Real-Time Data Capture (oscilloscope-type function)

Any control or feedback capability provided to the user interface must be implemented using the WBAN communication protocol. The protocol provides the tools enabling control of the WBAN and defines what can and cannot be accomplished. We strived to keep a simple set of WBAN message types and still implement complex user interface functions and application flexibility.

Sensor Node Identification, Association and Calibration

Sensor node identification requires a method for uniquely identifying a single sensor node to associate the node with a specific function during a health monitoring session. For example, a motion sensor placed on the arm performs an entirely different function than a motion sensor placed on the leg. Because two motion sensors are otherwise indistinguishable, it is necessary to identify which sensor should function as an arm motion sensor and which sensor should function as a leg motion sensor. In order to make node identification user friendly and intuitive, we developed a scheme

taking advantage of the inherent motion sensing capabilities of each sensor. We let the user arbitrarily place a motion sensor on his arm or leg, and then we are able to identify and associate the sensor with the proper function through a series of easy to follow instructions. Our form instructs the user to “move the arm sensor” or in a denser WBAN, “move the left arm sensor”. This interface is more intuitive from a user’s perspective, but was implemented using only WBAN protocol event messages already implemented for event processing. While the user is moving the sensor, the PS broadcasts an *ACTIS_EVENT_MASKMSG* requesting all sensors to report activity level estimations. Based on the largest activity estimate returned, the PS can identify which sensor the user is moving and associate it with the appropriate function.

The same message can also be used for configuration. Although the user interface presents these as two distinct functions, they are implemented using the same message. Event mask messages are also used to determine the degree of signal processing and the specific events of interest during a given health monitoring session. This approach allowed us to minimize the complexity of the communication protocol and still provide rich feature set to the application designer and user.

The Personal Server and ActiS nodes support two types of calibration. The first type is a sensor calibration; its purpose is to accommodate sensor-to-sensor variations and the exact nature of the calibration is sensor dependent. This is typically a one-time calibration and not expected to be a long-term function of the user interface, but certainly necessary for sensor preparation. The second type of calibration is a *session calibration*, required immediately prior starting a new monitoring session to calibrate the sensor in the context of its current environment. For example, Activity sensors on the leg might need an initial calibration of default orientation on the body.

Event Processing

The Personal Server is solely responsible for collecting data and events from the WBAN. Each sensor node in the network is sampling, collecting, and processing data. Depending on the type of sensor and the degree of processing specified at configuration, a variety of events will be reported to the Personal Server. An event log is created by aggregating event messages from all the sensors in the WBAN; the log must then be inserted into a session archive file. The Personal Server must recognize events as they are received and make decisions based on the severity of the event. Normally R-peak or heartbeat events do not create alerts, and are only logged in the event log. However, the Personal server will recognize when the corresponding heart rate exceeds predetermined threshold values. The Personal Server can alert the user that his heart rate has exceeded the target range. In addition to the regular status report in each super frame, the following events are currently supported:

- STEP (includes timestamp, step length, maximum forces)
- RPEAK – detection of heart beats using recognition of the R phase; the system generates a precise timestamp and time interval between the current and a previous heart beat or R-R interval (RRINT)
- Sensor Error (such as unexpected sensor reset)
- Force Threshold Exceeded - above normal accelerations (potential fall condition)
- User’s activity – AEE (one-second integration of the 3D motion vector)
- Triggered user’s activity – generates an event if the AEE exceeds a specified threshold.

Figure 11 illustrates a typical session and how WBAN protocol messages are used to calibrate sensors, start a session, and receive events.

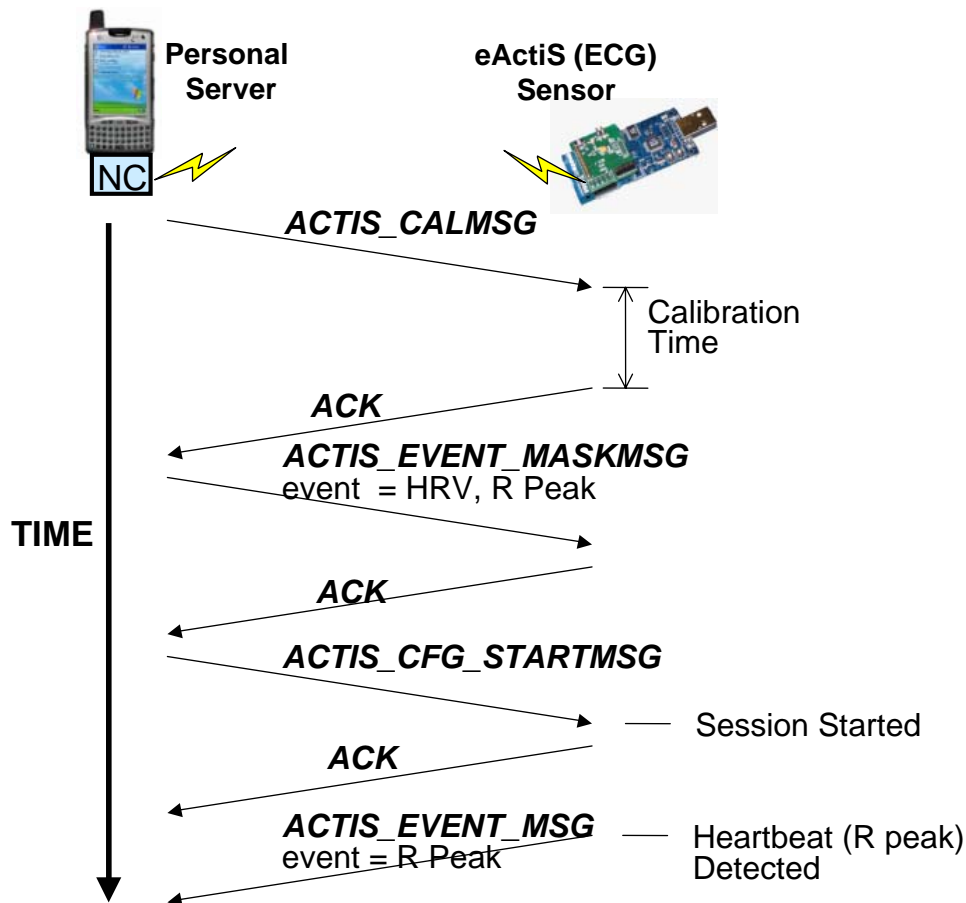


Figure 11 ActiS session initialization and event detection

Real-time Data Capture

Although all sensors in our system perform on-sensor processing and event detection, there are events where processed and summary events are not sufficient and real-time raw signal capture is necessary. During development, it was invaluable to be able to monitor sensor data in real-time. For heart rate sensors we implemented a single graphical ECG trace; for motion sensors we implemented three traces representing x, y, and z acceleration components on the same graph, as represented in Figure 9. This captured data is also stored to a file and can be analyzed off-line to improve step detection algorithms. In most cases, the algorithms were first developed on sample data sets previously recorded. When the algorithms worked well on the sample data sets, they were then implemented on the embedded sensors to run in real-time.

Even in a deployed system where intelligent sensors analyze raw data, process, and transmit application event messages, there may be cases where it is necessary to transmit raw physiological data samples. Such cases become apparent when considering a deployed ECG monitor. When embedded signal processing routines detect an arrhythmic event, the node should send an event message to the PS which will then be relayed to the appropriate medical server. The medical server, in turn, will provide an alert to the patient's physician. However, a missed heart beat can also be caused by electrode

movement. Therefore, it would be useful to augment this event with actual recording of the fragment of unprocessed ECG sensor data. The recording can be used by a physician to evaluate the type and exact nature of the event or to dismiss it as a recording artifact. In this case, the embedded sensor will begin streaming the real-time data to the personal server during a predefined time period. Figure 12 illustrates how a detected event can trigger a real-time data stream.

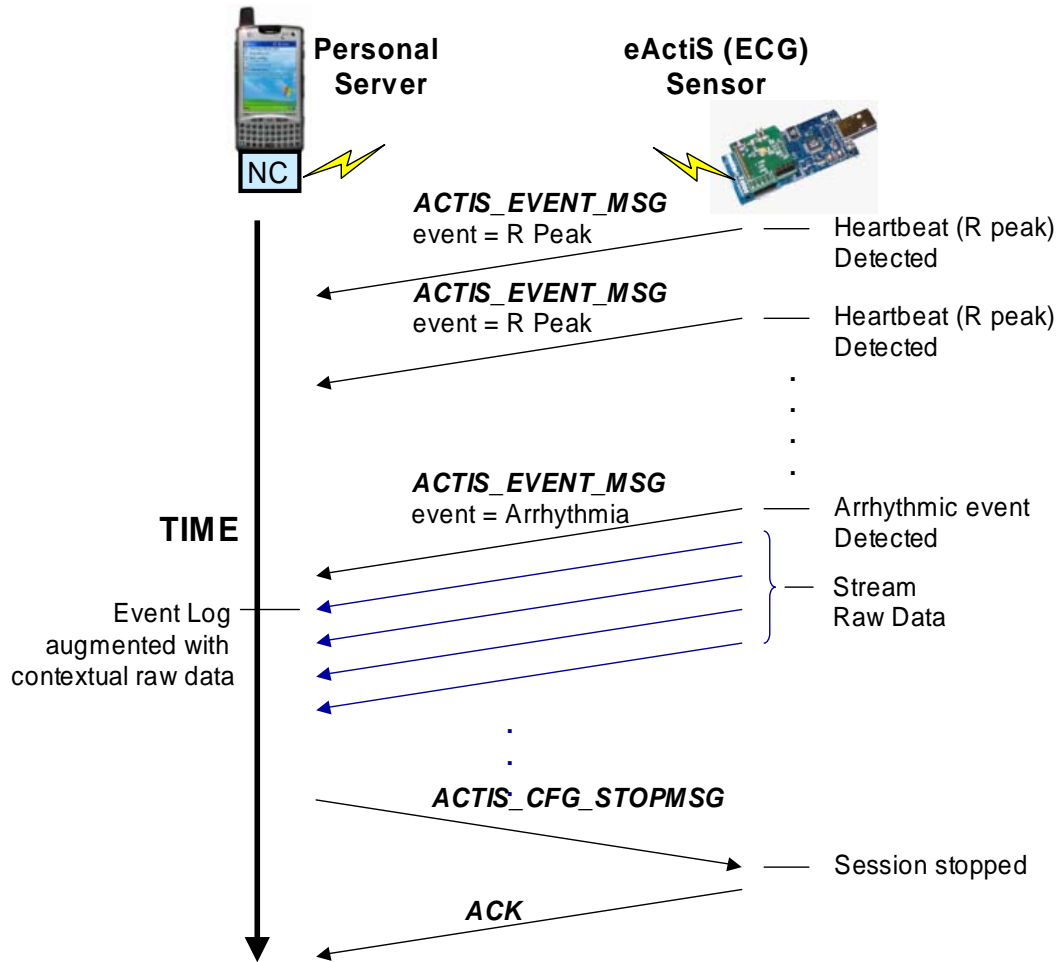


Figure 12 Arrhythmic event followed by real-time streaming data.

6 Conclusions

WBAN systems that monitor vital signs promise ubiquitous, yet affordable health monitoring. We believe that WBAN systems will allow a dramatic shift in the way people think about and manage their health – in the same fashion the Internet has changed the way people communicate to each other and search for information. This shift toward more proactive preventive healthcare will not only improve the quality of life, but will also reduce healthcare costs.

The proliferation of wireless devices and recent advances in miniature sensors prove the technical feasibility of a ubiquitous health monitoring system. However, WBAN designers face a number of challenges in an effort to improve user's compliance that depends on the ease of use, size, reliability, and security. In order to address some of these challenging tasks we have designed a WBAN prototype that includes accelerometer-based motion sensors, an ECG sensor, and a pocket PC based personal server. In this paper we describe both hardware and software architecture of our prototype. Our hardware architecture leverages off-the-shelf commodity sensor platforms. Similarly, our software architecture builds upon TinyOS, a widely used open-source operating system for embedded sensor networks. We are currently developing application of the prototype for computer assisted physical rehabilitation and health status monitoring.

Acknowledgements

The authors wish to thank Piet de Groen, Bruce Johnson, John Gober, and Reggie McMurtrey.

References

1. Bouten, C.V.C., Koekkoek, K.T.M., Verduin, M., Kodde, R., and Janssen, J.D., A Triaxial Accelerometer and Portable Data Processing Unit for the Assessment of Daily Physical Activity. in *IEEE Transactions On Biomedical Engineering*, 44 (3). March 1997. 136-147.
2. CardioLabs, <http://www.cardiolabs.com>, last accessed December 2005.
3. CardioNet, <http://www.cardionet.com>, last accessed December 2005.
4. Cox, D., Jovanov, E., and Milenkovic, A., Time Synchronization for ZigBee Networks. in *Proceedings of the 37th SSST*, (Tuskegee, U.S.A., 2005).
5. DeVaul, R.W., Sung, M., Gips, J., and Pentland, S., MITHril 2003: Applications and Architecture. In *Proceedings of ISWC 2003*, (White Plains, U.S.A., 2003).
6. Dishman, E., Inventing Wellness Systems for Aging in Place. in *IEEE Computer*, 37 (5). May 2004. 34-41.
7. Istepanian, R.S.H., Jovanov, E., and Zhang, Y.T., Guest Editorial Introduction to the Special Section on M-Health: Beyond Seamless Mobility and Global Wireless Health-Care Connectivity. in *IEEE Transactions on Information Technology in Biomedicine*, 8(4). December 2004. 405 - 414.
8. Jovanov, E., Milenkovic, A., Otto, C., and de Groen, P.C., A Wireless Body Area Network of Intelligent Motion Sensors for Computer Assisted Physical Rehabilitation. in *Journal of NeuroEngineering and Rehabilitation*, 2 (6). March 2005.
9. Lorincz, K., Malan, D.J., Fulford-Jones, T.R.F., Nawoj, A., Clavel, A., Shnayder, V., Mainland, G., Moulton, S. and Welsh, M., Sensor Networks for Emergency Response: Challenges and Opportunities. in *IEEE Pervasive Computing, Special Issue on Pervasive Computing for First Response*, 3 (4). October 2004. 16-23.
10. Maglaveras, N., Stamkopoulos, T., Pappas, C., Strintzis, M.G., An Adaptive Backpropagation Neural Network for Realtime Ischemia Episodes Detection: Development and Performance Analysis Using the European ST-T Database. in *IEEE Transactions on Biomedical Engineering*, 45 (7). July 1998. 805-813.
11. Maróti, M., Kusy, B., Simon, G., and Lédeczi, Á., The Flooding Time Synchronization Protocol. in *Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems*, (Baltimore, U.S.A., 2004).
12. Martin, T., Jovanov, E., and Raskovic, D., Issues in Wearable Computing for Medical Monitoring Applications: A Case Study of a Wearable ECG Monitoring Device. in *Proceedings of ISWC 2000*, (Atlanta, U.S.A., 2000).

13. Milenkovic, A., Milenkovic, M., Jovanov, E., and Hite, D., An Environment for Runtime Power Monitoring of Wireless Sensor Network Platforms. in Proceedings of the 37th SSST, (Tuskegee, U.S.A., 2005).
14. Moteiv, <http://www.moteiv.com>, last accessed December 2005.
15. National Coalition on Health Care. in <http://www.nchc.org/facts/cost.shtml>, accessed in August 2005.
16. Otto, C., Gober, J.P., McMurtrey, R.W., Milenković, A., and Jovanov, E., An Implementation of Hierarchical Signal Processing on Wireless Sensor in TinyOS Environment. in Proceedings of the 43rd ACMSE (Kennesaw, U.S.A., 2005).
17. Pan, J., Tompkins, W.J., A Real Time QRS Detection Algorithm. in IEEE Transactions on Biomedical Engineering, 32 (3). March 1985. 230-236.
18. Pentland, S., Healthwear: Medical Technology Becomes Wearable. in IEEE Computer, 37(5). May 2004. 34-41. 2004.
19. Shnayder, V., Chen, B., Lorincz, K., Fulford-Jones, T.R.F., and Welsch, M. Sensor Networks for Medical Care. in Harvard University Technical Report TR-08-05, 2005.
20. TinyOS, <http://www.tinyos.net>, last accessed December 2005.
21. U.S. Census Bureau, U.S. Interim Projections by Age, Sex, Race, and Hispanic Origin, <http://www.census.gov/ipc/www/usinterimproj/>, accessed in September 2005.
22. Warren, S., Lebak, J., Yao, J., Creekmore, J., Milenkovic, A., and Jovanov, E., Interoperability and Security in Wireless Body Area Network Infrastructures. in Proceedings of the 27th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, (Shanghai, China, 2005).
23. Wheelock, B., Autonomous Real-Time Detection of Silent Ischemia. M.S. thesis, University of Alabama in Huntsville, 1999.
24. ZigBee Alliance, <http://www.zigbee.org>, last accessed December 2005.