# DexterNet: An Open Platform for Heterogeneous Body Sensor Networks and Its Applications

*Philip Kuryloski*
*Annarita Giani*
*Roberta Giannantonio*
*Katherine Gilani*
*Raffaele Gravina*
*Ville-Pekka Seppa*
*Edmund Seto*
*Victor Shia*
*Curtis Wang*
*Posu Yan*
*Allen Yang*
*Jari Hyttinen*
*S. Shankar Sastry*
*Stephen Wicker*

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2008-174
http://www.eecs.berkeley.edu/Pubs/TechRpts/2008/EECS-2008-174.html

December 19, 2008

Acknowledgement

This work was supported in part by TRUST (The Team for Research in Ubiquitous Secure Technology), which receives support from the National Science Foundation (NSF award number CCF-0424422) and the following organizations: AFOSR (#FA9550-06-1-0244), Cisco, British Telecom, ESCHER, HP, IBM, iCAST, Intel, Microsoft, ORNL, Pirelli, Qualcomm, Sun, Symantec, Telecom Italia and United Technologies. This work was also supported in part by ARO MURI W911NF-06-1-0076, the Center for Information Technology Research in the Interest of Society (CITRIS), and Finnish Funding Agency for Technology and Innovation (Tekes).

# DexterNet: An Open Platform for Heterogeneous Body Sensor Networks and Its Applications

Philip Kuryloski ◇,†,○, Annarita Giani †, Roberta Giannantonio △, Katherine Gilani ⋆,
Raffaele Gravina ○, Ville-Pekka Seppä □, Edmund Seto ▽, Victor Shia †, Curtis Wang †,
Posu Yan †, Allen Y. Yang †, Jari Hyttinen □, Shankar Sastry †, Stephen Wicker ◇, Ruzena Bajcsy †
† Department of EECS, University of California, Berkeley, CA 94720
◇ Department of ECE, Cornell University, Ithaca, NY 14853
▽ School of Public Health, University of California, Berkeley, CA 94720
△ Telecom Italia, Turin, Italy
○ WSN Lab sponsored by Pirelli and Telecom Italia, Berkeley, CA 94704
□ Department of Biomedical Engineering, Tampere University of Technology, Tampere, Finland
⋆ Department of EE, University of Texas at Dallas, TX 75080.

**Abstract**

In this paper, we present an open-source platform for wireless body sensor networks called DexterNet. The system is motivated by shifting research paradigms to support real-time, persistent human monitoring in both indoor and outdoor environments. The platform utilizes a three-layer architecture to control heterogeneous body sensors. The first layer, called the body sensor layer (BSL), deals with design of different wireless body sensors and their instrumentation on the body. We detail two custom-built body sensors: one measuring body motions and the other measuring the ECG and respiratory patterns. At the second layer, called the personal network layer (PNL), the wireless body sensors on a single subject communicate with a mobile base station, which supports Linux OS and the IEEE 802.15.4 protocol. The BSL and PNL functions are abstracted and implemented as an open-source software library, called Signal Processing In Node Environment (SPINE). A DexterNet network is scalable, and can be reconfigured on-the-fly via SPINE. At the third layer, called the global network layer (GNL), multiple PNLs communicate with a remote Internet server to permanently log the sensor data and support higher-level applications. We demonstrate the versatility of the DexterNet platform via three applications: avatar visualization, human activity recognition, and integration of DexterNet with global positioning sensors and air pollution sensors for asthma studies.

## I. Introduction

Wireless body sensor networks (BSNs) have been an emerging research area in the field of sensor networks in the past five years. The development is mainly due to two reasons: 1. Continuing progress in the integration and miniaturization of sensors, processors, and radio devices. 2. Rising demand for advanced body sensor systems from pivotal areas of elderly protection and clinical patient monitoring to much broader applications in military, preventive healthcare, and consumer electronics. Traditional BSNs mainly involve single wearable sensors, such as fall detection [5], [7], [16], [18], walk and gait-phase detection [3], [13], and pulse-oximetry monitoring [11], [12]. More sophisticated systems consist of multiple heterogeneous sensors, adopt certain hierarchical architectures for real-time sensor management, and may integrate body sensors with other environmental sensors. Some examples include CodeBlue [9], MobiCare [4], and ALARM-NET [19]. These systems instrument the human body as an active mobile platform, and have the necessary mobility to support persistent monitoring in people's normal living environments.

In this paper, we present a novel platform for heterogeneous body sensor networks called *DexterNet*. The design principles of DexterNet are manifold:

1) DexterNet supports an open-source on-node signal processing library, namely, SPINE (Signal Processing In Node Environment) [17]. To our best knowledge, SPINE is the only open-source library that is versatile enough to support heterogeneous body sensors. We have designed and manufactured two different body sensors: one measuring body motions and the other measuring the ECG and respiratory patterns. The system can also conveniently integrate other commercially available sensor nodes, such as SHIMMER and MICAz. As a result, higher-level applications using DexterNet can seamlessly control different types of body sensors via the SPINE library.

This work was supported in part by TRUST (The Team for Research in Ubiquitous Secure Technology), which receives support from the National Science Foundation (NSF award number CCF-0424422) and the following organizations: AFOSR (#FA9550-06-1-0244), Cisco, British Telecom, ESCHER, HP, IBM, iCAST, Intel, Microsoft, ORNL, Pirelli, Qualcomm, Sun, Symantec, Telecom Italia and United Technologies. This work was also supported in part by ARO MURI W911NF-06-1-0076, the Center for Information Technology Research in the Interest of Society (CITRIS), and Finnish Funding Agency for Technology and Innovation (Tekes). Corresponding author: P. Kuryloski (pjk25@cornell.edu).

2) Harnessing the rich functionalities in SPINE, DexterNet supports real-time signal collection and sensor management. The system can be dynamically configured over the air. It provides a set of on-node services that can be tuned and activated by the user depending on different application needs.
3) To support long-term monitoring of multiple human subjects in both indoor and outdoor environments, DexterNet adopts a flexible three-layer BSN architecture. A body sensor layer (BSL) deals with the design of different sensors and their instrumentation on the body. A personal network layer (PNL) manages communication between the wireless body sensors and a mobile computer station. The mobile station can be either a computer or a smart phone that supports Linux OS and the IEEE 802.15.4 protocol. Finally, a global network layer (GNL) via the Internet permanently logs the sensor data and supports other higher-level applications on one or more secured network servers.

Figure 1 shows the three-layer architecture of DexterNet. At the BSL, the system supports two types of custom-built wireless wearable sensors. The first is a motion sensor board that consists of a triaxial accelerometer and biaxial gyroscope. The second is a biological sensor (biosensor) called *Wisepla* [14], which integrates an electrical impedance pneumography (EIP), an electrocardiogram (ECG), and a triaxial accerometer. Each sensor board then connects with a sensor network mote to form a wearable sensor mote. Here we choose the commercially available TelosB board. At the PNL, the body sensors communicate with a Nokia N800 series Internet tablet via a TelosB base-station board. The SPINE functions installed on the body sensors and the N800 manage the data collection, processing, and transmission of the data, and can be controlled via commands issued from the N800.

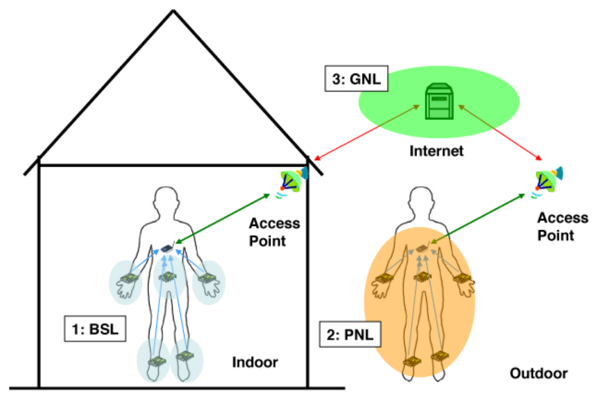


Fig. 1. The three-layer architecture of the DexterNet system: 1. Body sensor layer (BSL). 2. Personal network layer (PNL). 3. Global network layer (GNL). The pivotal component of the system is a Nokia N800 tablet at the PNL that communicates both to the BSL via IEEE 802.15.4 and the GNL via other broadband wireless channels.

Equipped with the versatile three-layer architecture and the open-source on-node library SPINE, DexterNet presents a competitive framework to support a variety of applications in healthcare, military, and consumer electronics. For example, a fall detection function has been implemented at the BSL level using SPINE on-node functions. In our implementation, each motion sensor is capable of outputting a binary decision of a falling event. Such functions reduce the amount of data that has to be transmitted between the nodes and the base station. More sophisticated applications such as human activity recognition and reconstruction of a graphical avatar for 3-D visualization can be implemented at the PNL level, which rely on the full-body motion data measured by multiple motion sensors at different key locations of the body (as shown in Figure 2).

### *A. Related Work*

Similar to DexterNet, many existing BSN platforms embrace a hierarchical architecture for real-time sensor control and data management. Some representative platforms are shown in Table I. A more comprehensive literature overview can be found in [10], [19], [22].

HealthGear [12] is a single-modality sensor network that integrates a low-power pulse oximeter with a smart phone via Bluetooth. CodeBlue [9] is a wireless sensor platform intended for deployment in emergency medical care. It integrates a pulse oximeter and a ECG sensor with PDAs and PCs to enhance seamless transfer of data among caregivers. The platform uses IEEE 802.15.4 as the wireless protocol, and is intended to scale in dense networks with volatile network conditions.

WWBAN [10] adopts a three-layer multi-sensor platform that is similar to DexterNet. Multiple motion sensors and ECG sensors are placed on the human body. They communicate with either a PDA or a PC to provide a transparent interface to the user, and an interface to the (remote) medical server using the Internet. However, the system is mainly comprised of proprietary software, and it does not provide an open-source library such as SPINE to support on-node computation and decision-making.

Finally, ALARM-NET [19] belongs to a group of wireless sensor networks for assisted living. The focus of the system is the integration of body sensors with environmental sensor networks in a scalable and heterogeneous architecture. ALARM-NET

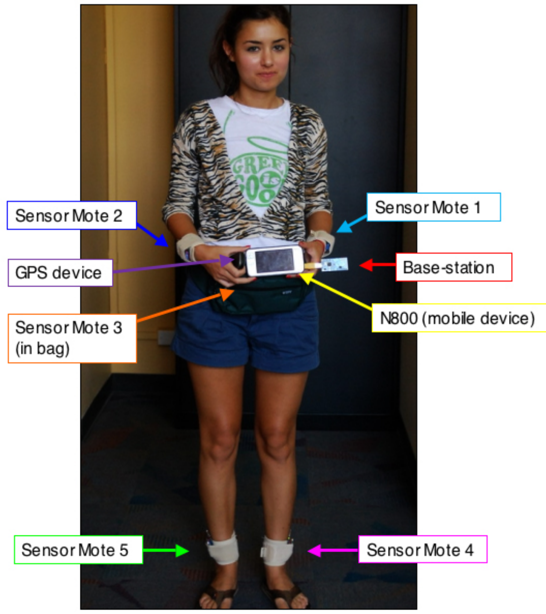


Fig. 2. Illustration of the DexterNet system instrumented on a wearer. The deployment includes five motion sensor motes, a Nokia N800 tablet, and a GPS positioning sensor.

TABLE I
COMPARISON OF EXISTING BODY SENSOR NETWORKS WITH THE DEXTERNET PLATFORM.

| Platforms | Sensor Devices | Base Devices | Node Protocols | Open Source | Environmental Sensors |
|---|---|---|---|---|---|
| HealthGear [12] | pulse oximeter | smart phone | Bluetooth | No | No |
| CodeBlue [9] | pulse oximeter<br>ECG, SHIMMER | PC<br>PDA | 802.15.4 | No | No |
| WWBAN [10] | motion, ECG | PC, PDA | 802.15.4 | No | No |
| ALARM-NET [19] | pulse oximetry<br>motion, ECG | STARGATE<br>PDA, PC | Bluetooth<br>802.11 | No | Yes<br>(temperature, light, PIR) |
| DexterNet | motion, ECG<br>EIP, GPS<br>MICAz, SHIMMER | PDA<br>PC | 802.15.4 | Yes | Possible via SPINE<br>(e.g., air pollution sensor) |

uses MICAz sensors and STARGATE to relay the information from body sensors and environmental sensors to PDAs and PCs in a large and complex indoor setting using either Bluetooth or the 802.11 protocol.

The rest of the paper is organized as follows: Section II proposes the overall architecture of DexterNet and explains the relationship among different components of the three-layer hierarchy. Based on the hierarchy, Section III first discusses the design and specification of the body sensors in the bottom layer BSL. Section IV then discusses the open-source SPINE network that provides software services and control of both BSL and PNL. Section V showcases three high-level applications: 1. Avatar visualization. 2. Human activity recognition. 3. Integration of DexterNet with portable air pollution sensors for the study of asthma attack. Finally, Section VI discusses some limitations of the current implementation and future directions.

## II. SYSTEM ARCHITECTURE

DexterNet provides a rich, complete pathway for sensing, distributed processing, wireless communication, and data fusion, which serves as a foundation for higher-level applications. Although subsets of these functionalities have been implemented in other systems, often these sets are overlapping. We hope that in providing an open platform with DexterNet, variations in functionality can be built using a common base. This will avoid redundant development efforts in different sensor network systems. Additionally, the diverse nature of our team has driven the design requirement for DexterNet to provide maximum flexibility and extensibility, with maximum potential for reusability of its components.

The structure of DexterNet is shown in Figure 3. The open-source SPINE framework provides the flexibility in constructing physical components of the system at the BSL and PNL layers. Particularly, SPINE has been developed such that there is separation in code of its sensing, processing, and communication functions. As a result, SPINE is portable across TinyOS mote platforms, and easily extends to support new sensors through the use of sensor drivers.

All SPINE functionalities are dynamically configured over the air, helping to achieve the runtime flexibility and reconfigurability at the BSL and PNL layers. SPINE includes a base station component, and allows the use of a Nokia N800 or

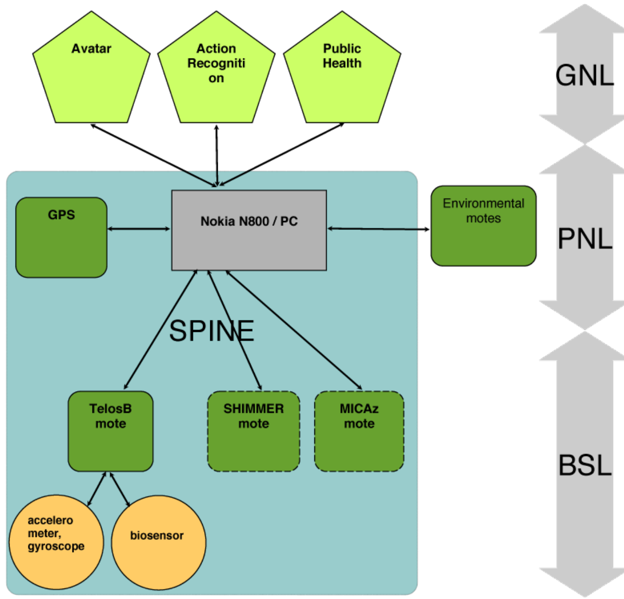


Fig. 3. Architecture of the DexterNet system. The Body Sensing Layer (BSL) includes motes and attached sensors. The Personal Network Layer (PNL) includes the N800 portable base station and associated sensors. The BSL and PNL are driven by SPINE. The Global Network Layer (GNL) includes our applications built with the DexterNet system.

PC at the PNL layer. The N800 allows the wearable system to be portable, and allows for the integration of GPS and other environmental sensors. Our experience has shown that integration of various commercial devices, such as the N800, has come with considerable effort. DexterNet aims to maximize the utility of such efforts.

Each of our applications, including avatar visualization, action recognition and asthma studies, is built on top of the SPINE base station API. They each configure the appropriate sensors and on-node signal processing according to their specific goals and requirements. These applications need not depend on any specific sensor program or directly interact with the BSL. This ensures that all applications benefit from enhancements made at the BSL and PNL of the system, such as improvements in robustness, capacity, and energy consumption. Furthermore, developers can simultaneously work on application-level software as well as SPINE software without the tight coupling required in traditional application-specific sensor network systems.

## III. Design of Body Sensors

### A. *Motion Sensors*

DexterNet supports the deployment of multiple motion sensor nodes placed at different body locations (see Figure 2), which communicate with a base station. The sensor nodes and the base station are built using the TelosB boards. TelosB runs TinyOS on an 8 MHz microcontroller with 10K RAM and communicates using the 802.15.4 wireless protocol. Each custom-built sensor board has a triaxial accelerometer and a biaxial gyroscope, which is attached to TelosB (shown in Figure 4). Each axis is reported as a 12bit value to the node, indicating values in the range of $\pm 2g$ and $\pm 500^\circ/s$ for the accelerometer and gyroscope, respectively. The battery life of continuous measurement and wireless raw data output is approximately 20 hours.

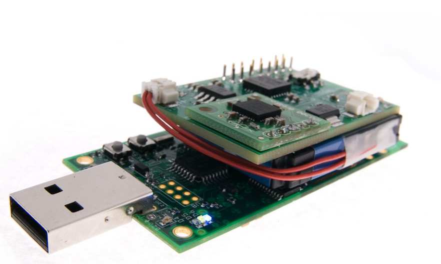

Fig. 4. Illustration of the motion sensor node. The sensor board on the top is a custom-built motion sensor with a triaxial accelerometer and a biaxial gyroscope. The middle layer is a Li-ion battery. The sensor board on the bottom is a standard TelosB network node.

The current hardware design of the sensor contributes certain amounts of measurement error. The accelerometers typically require some calibration in the form of a linear correction, as sensor output under $1g$ may be shifted up to 15% in some sensors. It is also worth noting that the gyroscopes produce an indication of rotation under straight-line motions. Fortunately, these systematic errors appear to be consistent across experiments for a given sensor board. However, without calibration to correct them, the errors may affect the action recognition if different sets of sensors are used interchangeably in the experiment. [1]

### *B. Biosensors*

The biosensor is capable of measuring acceleration, electrocardiogram (ECG), and electrical impedance pneumography (EIP) through four small electrodes connected to the side of the ribcage of the subject, as shown in Figure 5 and 6. The ECG signal is used to derive heart rate and heart rate variability (HRV). The EIP signal is produced by respiration and can be used to derive a variety of breathing related parameters like respiration rate, minute ventilation volume, flow/volume curve, and inspiration/expiration times.

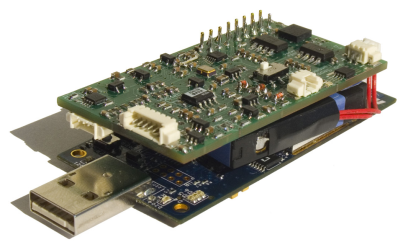

Fig. 5. Illustration of the biosensor board (top) connected to the TelosB network node (bottom). The middle layer is a Li-ion battery. The white connector on the left side is for the four skin electrodes.

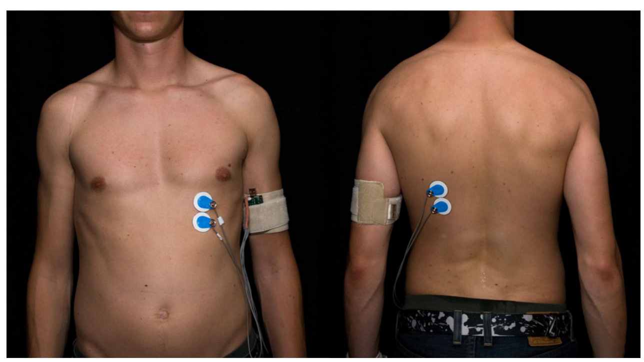

Fig. 6. The electrode locations are chosen to obtain a good signal-to-artefact ratio (SAR) in both EIP and ECG signals. Electrodes in both pairs are placed vertically next to each other. The front-end pair is placed vertically right below the *pectoralis major* and horizontally in the middle between the side of the body and mid axillary line. The back-end pair is placed on the same vertical and horizontal location to create a sensitivity field through and around the left lung.

Single channel ECG measurement is quite straightforward. The main challenge is breathing measurement with EIP technique especially for volumic parameters. So far we have tested the accuracy of the system during ergometer and running exercises [14], [15]. The accuracy of breathing minute volume assessment was degraded due to intense motion of the body during running, but results with average relative error of 11% were still obtained. Also the effect of different electrode placements on movement error susceptibility has been studied [8].

The biosensor runs real time signal processing algorithms that detect events in heart and respiration signals and calculate physiological parameters from them. The requested parameters are sent to the base station through the SPINE framework. This reduces the amount of network traffic compared with raw data sending and enables using a higher sampling rate needed in accurate HRV analysis. The possibility to assess breathing-related parameters separates this sensor from most of the similar projects. Cardiac and pulmonary measurement together provide data that can be used to derive high-level physiological parameters related to physical and mental state of the subject.

The biosensor is connected to a TelosB and is in the same form factor as the motion sensor. The battery life of continuous measurement and raw data output is approximately 20 hours, similar to the motion sensor.

[1] More sophisticated motion sensors do exist in the industry, which can utilize heterogeneous sensor fusion techniques to self-calibrate the accelerometer and gyroscope. One example is the Microstrain Gyro Enhanced Orientation Sensor at: http://www.microstrain.com/.

### C. Other Compliant Sensors

The heterogenity of DexterNet allows a wide variety of sensors and motes to be integrated into the system. The SPINE framework provides support for the Intel SHIMMER and the MICAz motes as well as any sensors that can be attached. The SHIMMER has an onboard accelerometer, MicroSD slot, and ADC converters for attaching external sensors. The MICAz has many sensors available as addons, including sensors such as GPS, humidity, barometric pressure, ambient light, sound, magnetometer, etc.

Our current system has a Bluetooth GPS sensor that directly interfaces with SPINE. Since the GPS unit itself is not a SPINE node, the data integration is done at a different layer than SPINE nodes such as the motion sensor and biosensor. The GPS provides longitude and latitude coordinates primarily in outdoor environments at a speed of 1 Hz.

## IV. The SPINE Framework

SPINE (Signal Processing In Node Environment) [2] is an open-source framework for distributed signal processing algorithms in wireless sensor networks (WSNs). The functional architecture of SPINE is shown in Figure 7. It provides a set of on-node services that can be tuned and activated by the user depending on different application needs. The open-source framework speeds up the design of WSN applications through high-level abstractions and provides support to quickly explore implementation tradeoffs through fast prototyping. SPINE also provides an efficient wireless communication protocol for dynamic network configuration and management.

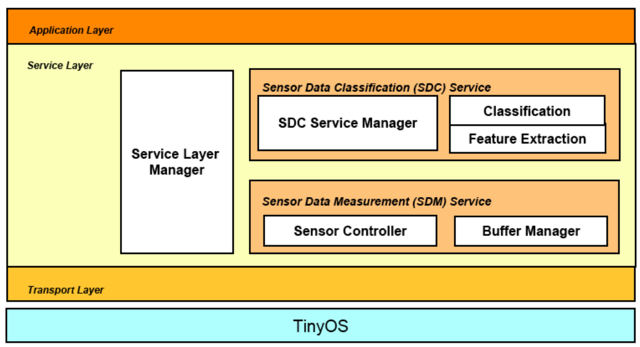


Fig. 7. The SPINE functional architecture.

The SPINE framework has two main modules, one for the sensor node side, and the other for the server/base-station side. The node module is developed in TinyOS 2.x environment. It provides the following three on-node service components: 1. Communication. 2. Sensing. 3. Signal processing. Accordingly, the source code of the module is organized in a similar manner. The communication component utilizes a *time division multiple access* (TDMA) protocol to avoid packet collision. In addition, a simple power saving mechanism can be activated where nodes listen for messages only briefly after sending a message. All sensor drivers implemented in the sensing component appear similar to the signal processing and communication components. As a result, any new sensor driver will be immediately available for all processing components. The signal processing component provides two different modules: the first one periodically performs feature extraction on sensor data and reports it, and the second one reports chosen features conditionally based on some thresholds. These design features make it easy to extend the SPINE framework, and allow various team members to develop different parts of the framework simultaneously. SPINE sensing and processing functionalities are dynamically configured through over-the-air messaging. This allows each application supported by the system to reconfigure the SPINE network as desired, quickly and easily.

The server module is implemented in Java SE and acts as the coordinator of a sensor network. It consists of functionalities that activate and control on-node services depending on the application requirements. The implementation instead does not use any TinyOS specific APIs and can be run independently on the underlying protocol stack (e.g., the ZigBee network). This has allowed the use of a Nokia N800 tablet as a handheld base station for the wearable sensor network. The N800 provides a platform for Bluetooth and Wi-Fi connectivity to allow forwarding of data to the GNL. It allows the realization of a body sensor network that can operate both inside the home and outdoor, a key feature for supporting a wide variety of human monitoring applications.

## V. Applications and Evaluation

### A. Avatar

We first demonstrate an application called *Avatar*, which uses a network of motion sensors on the human body to reconstruct and visualize the wearer's full-body motion in real time. The application can be used to remotely monitor and assess the well

[2]The SPINE software is available for download at http://spine.tilab.com/.

being of elderly people living alone. It can also be used in tele-healthcare for physicians to remotely record and visualize the movements of patients. Avatar provides much of the same information about activity that could be captured by video, but does so providing a considerably higher level of privacy for the monitored person. This is quite important because it is unlikely that the average person would be willing to accept continuous video surveillance of their home. Additionally, Avatar has the benefit of being derived from wearable sensors, and so is portable. For the purpose of visualizing motion, a configuration of five nodes (one on each leg, one on each arm and one on the torso) are the minimum number of sensors required, which will be used in this section. To provide finer measurement of the full-body movement, more sensor nodes can be worn by a person.

Avatar makes use of the Java Monkey Engine (jME) [1] and physics plug-in [2] to render and animate a graphical human avatar. jME allows us to create an underlying skeleton with joints, then use sensor data to continuously change the orientation of this skeleton.

Through SPINE, each node estimates the pitch and roll of its orientation in space and reports this pair of values to the base station. The orientation in space of a single sensor node is computed based on the apparent direction of gravity as seen by the sensor board's accelerometer. When considered as a vector, the accelerometer will read the vector sum of gravity and acceleration resulting in movement of the sensor board. Under relaxed motion, motion component of the vector is less than 10% of the magnitude of the gravity vector. As a result, this motion component is neglected and we continuously interpret the direction of the accelerometer vector as the direction of gravity. Although the sensor board's gyroscope would presumably be beneficial in separating the gravity and motion vector components, in practice the gyroscopes indicate rotation even while the sensor is moved in a strictly translational (and non-rotational) path. As a result, the gyroscopes are currently ignored. This method is stateless and does not accumulate error, as would occur if accelerometer or gyroscope data were integrated to estimate velocity and displacement.

A snapshot of the output from Avatar is shown in Figure 8 with the physics skeleton in view. The yellow bars indicate the axes of the various skeleton joints, with gravity vectors are shown in red. At every frame, the orientation of the skeleton is compared to the data from the sensor nodes. A simulated force is then applied to each sensed body part to push it toward the orientation reported by the sensor. The force is such that the physics skeleton tracks the motion of the wearer, but is limited by the joints of the skeleton.

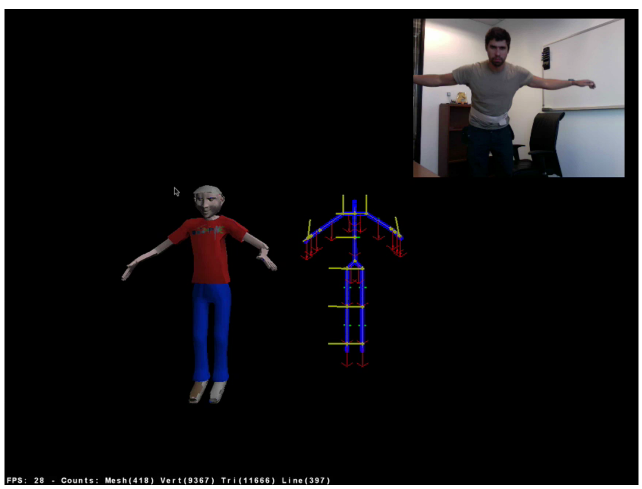


Fig. 8. A screen shot for the Avatar with overlaid image of the wearer.

### *B. Action Recognition*

In addition to using graphical avatars to visualize and analyze human poses and movements, another application of DexterNet is human action/activity recognition. Traditionally, human action recognition has been extensively studied in computer vision using camera sensors placed in an (indoor) environment where human users reside. Compared with these high-power, high-bandwidth camera systems, body sensor networks such as DexterNet have several distinct advantages: 1. Body sensor systems do not require instrumenting the environment with cameras or other sensors. 2. Body sensor systems also have the necessary mobility to support persistent monitoring of a subject during her daily activities in both indoor and outdoor environments. 3. With the continuing miniaturization and integration of mobile processors and wireless sensors, it has become possible to manufacture body sensor systems that can densely cover the human body to record and analyze very small movements (e.g., breathing and spine movements) with higher accuracy than most extant vision systems. Such action recognition systems have been used in applications such as medical-care monitoring, athlete training, tele-immersion, and human-computer interaction. For a detailed survey of the literature, the reader is referred to [21].

We have constructed an open-source benchmark database for human action recognition using the DexterNet system called *Wearable Action Recognition Database* (WARD). The purpose of WARD is to offer a public and relatively stable data set for quantitative comparison of existing and future algorithms for human action recognition using body motion sensors. The database has been carefully constructed under the following conditions:

1) The database contains sufficient numbers of human subjects with a large range of age differences.
2) The designed action classes are general enough to cover most typical activities that a human subject is expected to perform in her daily life.
3) The locations of the wearable sensors are selected to be practical for full-fledged commercial systems.
4) The sampled action data contain sufficient variation, measurement noise, and outliers in order for existing and future algorithms to meaningfully examine and compare their performance.

The data are sampled from 7 female and 13 male subjects (in total 20 subjects) with ages ranging from 19 to 75. For more details about the data collection, please refer to the human subject protocol included in the WARD database. The database also includes a MATLAB program to visualize the action data measured from the five motion sensors (Figure 9).[3]

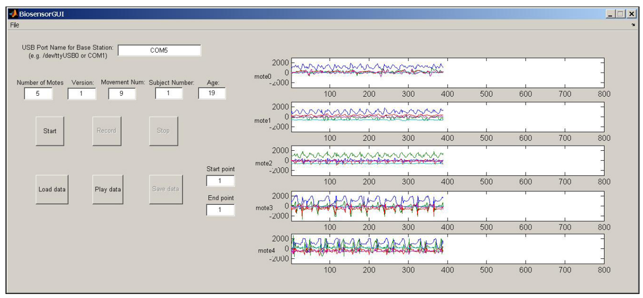


Fig. 9. A MATLAB program that interfaces with the TelosB base station via the series port. The program can receive, record, and replay accelerometer and gyroscope data from a network of motion sensors.

We have proposed a distributed recognition algorithm to classify human actions using the low-bandwidth motion sensors [20], [21]. These actions include transient actions, e.g., bending, lying down, and standing up; and continuous actions, e.g., walking, running, turning, and going upstairs. The algorithm classifies human actions using a set of training motion sequences as prior examples. It is also capable of rejecting outlying actions that are not in the training categories. The classification is operated in a distributed fashion on individual sensor nodes and a base station computer. More importantly, the algorithm is robust and adaptive to the change of active sensors in a body network on-the-fly due to either sensor failure or network congestion. The recognition precision only decreases gracefully using smaller subsets of active sensors. The accuracy of the framework is validated using the WARD database.

## *C. Public Health*

DexterNet has many applications within the field of public health, where the ability to objectively monitor the activity patterns of users may improve understanding of exposures to environmental hazards such as air pollution that are associated with asthma attacks, chronic obstructive pulmonary disease (COPD), cardiovascular disease, as well as premature mortality. The addition of the biosensor data provides a mechanism to monitor physiological responses to such exposures in real time that may be predictive of severe disease events (e.g., an asthma attack).

The inclusion of geographic location data from the GPS is also important for such applications. In the past, spatial epidemiologic studies have relied upon rather crude measures of location when describing a person's exposure to environmental hazards such as air pollution. For example, some studies have simply used residential locations as a proxy for a person's location. But in reality, individuals are mobile and have activity patterns that may include time away from home, at work, running chores, and exercising and playing. A system which allows for continual monitoring of an individual's location may greatly improve the assessment of exposures for such epidemiologic studies.

To evaluate the DexterNet system for such applications, we conducted a field experiment in which the system was used to collect and process an integrated set of data related to an individual's outdoor experience. The experiment consisted of a series of prescribed walks. A convenient sample of six adults (five male and one female) were asked to walk a 2.4 km route. The walk included sections that were uphill, downhill, and flat, as well as sections that were along a busy roadway, a downtown commercial/retail area, as well as a calmer path through a university campus. Over the course of the walk, various sensor data were logged, including triaxial accelerometry and biaxial gyroscopy (at the left wrist, waist, and left ankle positions), GPS location, and air pollution (airborne particulate matter $\leq 2.5\,$um in size, PM2.5). The motion data were logged at 30 Hz. GPS was logged at 1 Hz. The air pollution data were logged separately using a Met One Aerocet 531, a handheld particle counter

[3]The WARD database is available for download at: http://www.eecs.berkeley.edu/~yang/software/WAR/.

that takes 2-minute samples continuously during the walk. These data were combined and processed to ascertain specific information on the individual's experience (e.g., assessing the magnitude of physical activity in certain geographic locations, or the air pollution in each location that heavy activity occurred).

As an example, Figure 10 illustrates the GPS trace of the walking route. The application determines the changes in elevation during the walk from the GPS data. A motion sensor at the waist was used to derive energy expenditure using the Generalized Linear Model [6]. The breathing minute ventilation is derived from the biosensor EIP signal [15]. Heart rate is obtained from the biosensor ECG signal using a simple R-peak detection algorithm.

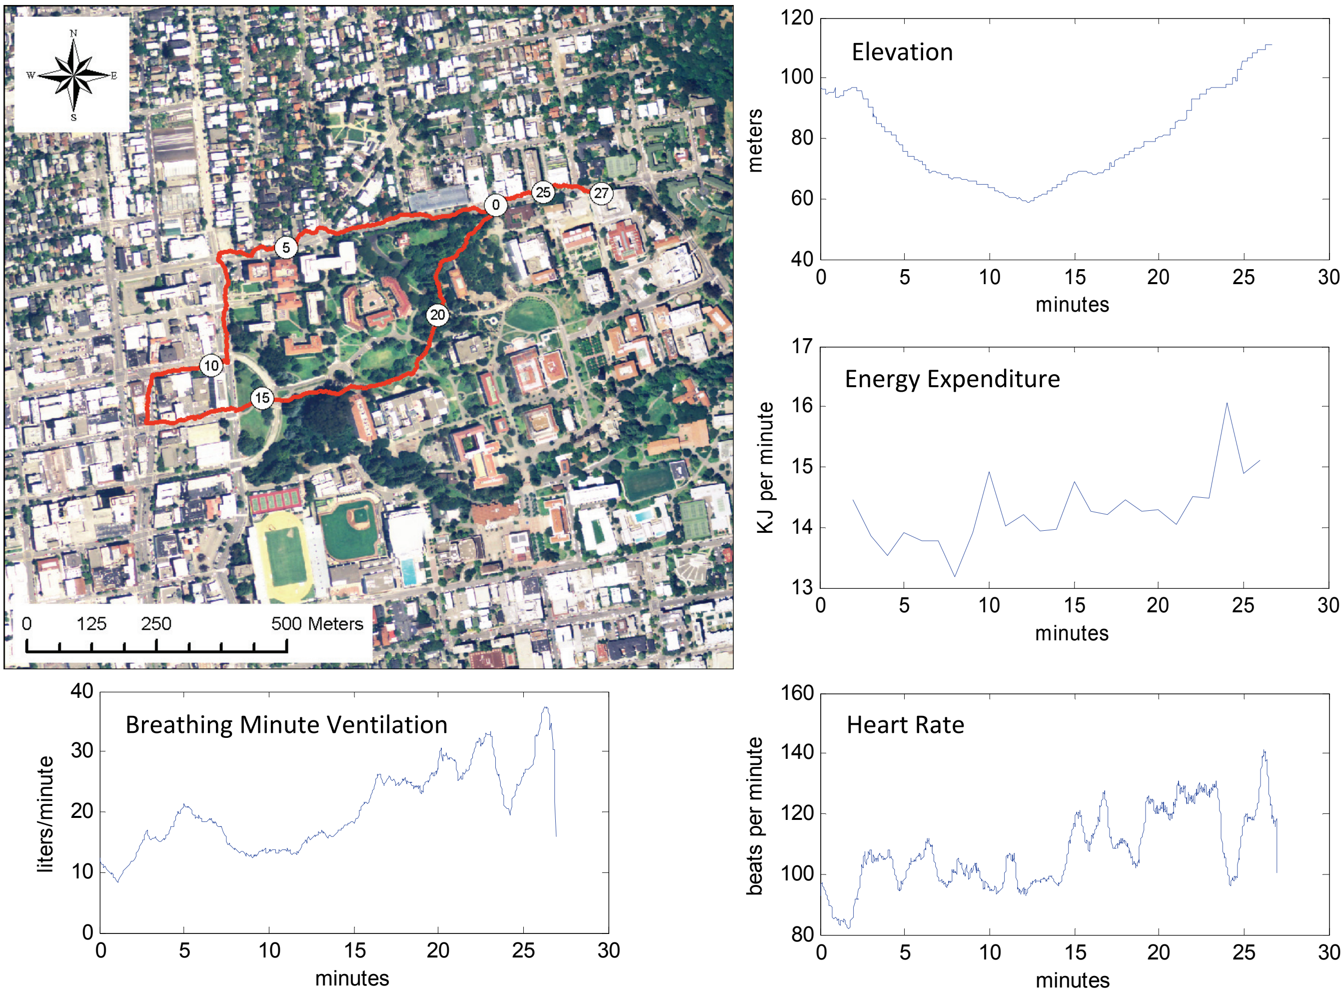


Fig. 10. GPS trace of campus walk with derived information from GPS, motion sensor and biosensor. Circles on the map indicate elapsed time in minutes.

The GPS data were also used to map PM2.5 concentrations from the Aerocet monitor during the walk (Figure 11). Data from three participants illustrates less spatial variability in pollutant concentrations than interday variability. We note that one of the days (the right panel of Figure 11) corresponds to a "Spare the Air Day", a day when an elevated air pollution warning (typically for ozone rather than PM2.5) was issued by the Bay Area Air Quality Management District to the general public. From these data it is possible to derive an individual's average and cumulative air pollution exposures and physiologic response for use in long-term epidemiologic studies.

## VI. Discussion and Future Directions

In this paper, we have discussed DexterNet, a novel platform for heterogeneous body sensor networks. The key attributes of DexterNet are twofold: 1. It promotes an open-source sensor environment that supports limited on-node computation, robust sensor communication, and online reconfigurable network management. 2. The platform is versatile enough to support a variety of existing body sensors and other future sensors that comply with the SPINE specifications. Through a hierarchy of three network layers, it resolves the dependency of higher-level applications toward the implementation of wireless body sensors and communication protocols.

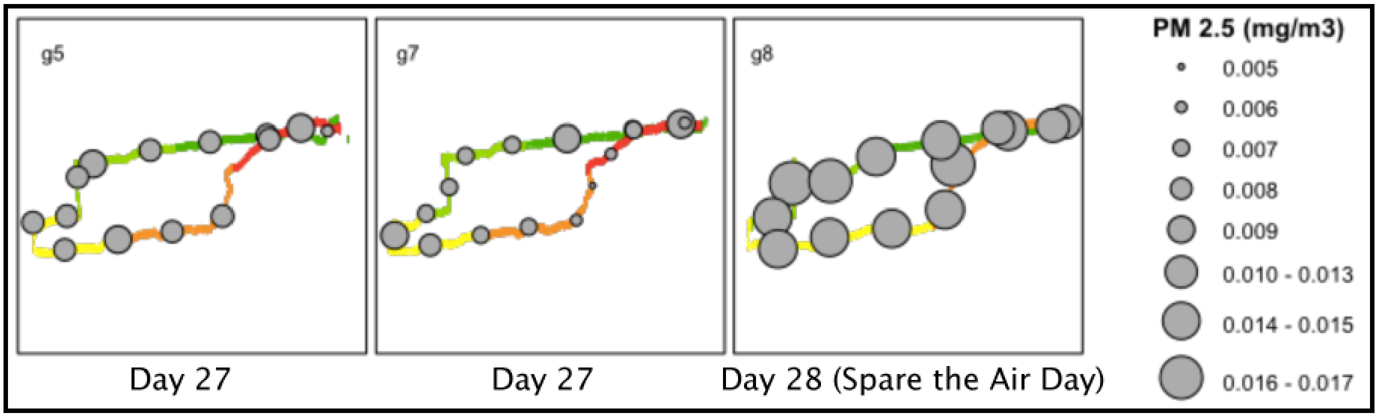


Fig. 11. GPS traces of campus walks for 3 participants, with geocoded PM2.5 measurements (circles), suggesting less spatial variability than interday variability. Note the right panel illustrates the data for a "Spare the Air Day", a day when an elevated air pollution warning was issued by the Bay Area Air Quality Management District to the general public.

One advantage of the DexterNet system is its low cost compared to other existing commercial systems that are more expensive and do not necessarily support open-source development. Currently, our system is limited by the choice of the off-the-shelf components (e.g., the TelosB mote and the N800), which in their current stages of development may not offer the most convenient form-factor and attractive packaging to make large-scale and long-term use practical. However, the limitation can be easily addressed by migrating the components to other commercial components. We are currently exploring new solutions to improve our system for future use.

There are numerous potential services that may be implemented through DexterNet, especially in the area of preventive healthcare. For example, It is possible through the classification algorithms described to identify conditions that are predictive of asthma attacks and warn users to reduce physical activity and/or move indoors. Such systems can also create maps of microscale air pollution when they are deployed in sufficiently large numbers. Currently, only regional air pollution maps are available from the sparsely located fixed-site monitoring that regulatory agencies implement.

An important consideration in the deployment of DexterNet is the need to protect the wearer's privacy, which is mandated by the 1996 Health Insurance Portability and Accountability Act (HIPAA). Work on private-key and public-key cryptography schemes for sensor networks is applicable, but must be integrated into an appropriate authentication and authorization framework. In addition, it is important to consider other security attacks, such as injection of anomalous data and illegal data exfiltration. Authentication, key establishment, robustness to denial-of-service attack, secure routing, and node capture are some of the security challenges in wireless sensor networks. In the case of BSN, these issues appear even more serious given the limited bandwidth, power supply, and storage and computational resources of the platform. When implementing privacy and security preserving features becomes critical to certain high-level applications, the use of a mobile computer station such as the N800 at the personal network layer reduces the burden on each wearable node. One can use the higher capacity of the base station to manage the privacy requirements across geographic locations and to authenticate various individuals, a scenario not necessarily possible with a multi-person BSL or a hybrid body and environmental sensor environment.

## Acknowledgments

The authors would like to thank Dr. Marco Sgroi at the WSN Lab Berkeley, Dr. Yuan Xue at the Vanderbilt University, and Dr. Roozbeh Jafari at the University of Texas, Dallas, for their valuable suggestions and literature references.

## References

[1] Java monkey engine (http://www.jmonkeyengine.com/), September 2008.
[2] jmephysics (https://jmephysics.dev.java.net/), September 2008.
[3] P. Barralon, N. Vuillerme, and N. Noury. Walk detection with a kinematic sensor: Frequency and wavelet comparison. In *Proceedings of the 28th IEEE EMBS Annual International Conference*, pages 1711–1714, 2006.
[4] R. Chakravorty. A programmable service architecture for mobile medical care. In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications Workshop*, 2006.
[5] J. Chen, K. Kwong, D. Chang, J. Luk, and R. Bajcsy. Wearable sensors for reliable fall detection. In *Proceedings of the IEEE Engineering in Medicine and Biology Conference*, pages 3551–3554, 2005.
[6] K. Chen and M. Sun. Improving energy expenditure estimation by using a triaxial accelerometer. *Journal of Applied Physiology*, 83:2112–2122, 1997.
[7] T. Degen, H. Jaeckel, M. Rufer, and S. Wyss. SPEEDY: A fall detector in a wrist watch. In *Proceedings of the IEEE International Symposium on Wearable Computers*, pages 184–187, 2003.
[8] O. Lahtinen, V-P. Seppä, J. Väisänen, and J. Hyttinen. Optimal electrode configurations for impedance pneumography during sports activities. In *Proceedings of the 4th European Congress for Medical and Biomedical Engineering*, 2008.
[9] D. Malan, T. Fulford-Jones, M. Welsh, and S. Moulton. CodeBlue: An ad hoc sensor network infrastructure for emergency medical care. In *Proceedings of the International Workshop on Wearable and Implantable Body Sensor Networks*, 2004.

[10] A. Milenkovic, C. Otto, and E. Jovanov. Wireless sensor networks for personal health monitoring: Issues and an implementation. *(in press) Computer Communications*, 2006.

[11] M. Morón, E. Casilari, R. Luque, and J. Gázquez. A wireless monitoring system for pulse-oximetry sensors. In *Proceedings of the 2005 Systems Communications*, 2005.

[12] N. Oliver and F. Flores-Mangas. HealthGear: A real-time wearable system for monitoring and analyzing physiological signals. In *Proceedings of the International Workshop on Wearable and Implantable Body Sensor Networks*, pages 61–64, 2006.

[13] I. Pappas, T. Keller, S. Mangold, M. Popovic, V. Dietz, and M. Morari. A reliable gyroscope-based gait-phase detection sensor embedded in a shoe insole. *IEEE Sensors Journal*, 4(2):268–274, 2004.

[14] V-P Seppä, J. Väisänen, P. Kauppinen, J. Malmivuo, and J. Hyttinen. Measuring respirational parameters with a wearable bioimpedance device. In *Proceedings of the 13th International Conference on Electrical Bioimpedance*, 2007.

[15] V-P. Seppä, J. Väisänen, O. Lahtinen, and J. Hyttinen. Assessment of breathing parameters during running with a wearable bioimpedance device. In *Proceedings of the 4th European Congress for Medical and Biomedical Engineering*, 2008.

[16] A. Sixsmith and N. Johnson. A smart sensor to detect the falls of the elderly. *Pervasive Computing*, pages 42–47, 2004.

[17] The SPINE Team. The spine manual version 1.2. Technical report, Telecom Italia Lab, 2008.

[18] G. Williams, K. Doughty, K. Cameron, and D. Bradley. A smart fall and activity monitor for telecare applications. In *Proceedings of the IEEE International Conference in Medicine and Biology Society*, 1998.

[19] A. Wood, G. Virone, T. Doan, Q. Cao, L. Selavo, Y. Wu, L. Fang, Z. He, S. Lin, and J. Stankovic. ALARM-NET: Wireless sensor networks for assisted-living and residential monitoring. Technical report, Department of Computer Science, University of Virginia, 2006.

[20] A. Yang, R. Jafari, P. Kuryloski, S. Iyengar, S. Sastry, and R. Bajcsy. Distributed segmentation and classification of human actions using a wearable sensor network. In *Proceedings of the CVPR Workshop on Human Communicative Behavior Analysis*, 2008.

[21] A. Yang, R. Jafari, S. Sastr, and R. Bajcsy. Distributed recognition of human actions using wearable motion sensor networks. *Submitted to Journal of Ambient Intelligence and Smart Environments*, 2008.

[22] J. Yick, B. Mukherjee, and D. Ghosal. Wireless sensor network survey. *Computer Networks*, 52(12):2292–2330, 2008.