

# Context-Aware Wireless Sensor Networks for Assisted-Living and Residential Monitoring

A. Wood, J. Stankovic, G. Virone, L. Selavo, Z. He, Q. Cao, T. Doan, Y. Wu, L. Fang, R. Stoleru

Department of Computer Science, University of Virginia

## Abstract

Improving the quality of healthcare and the prospects of “aging in place” using wireless sensor technology requires solving difficult problems in scale, energy management, data access, security, and privacy. We present AlarmNet, a novel system for assisted-living and residential monitoring that uses a two-way flow of data and analysis between the front and back-ends to enable context-aware protocols that are tailored to residents’ individual patterns of living.

AlarmNet integrates environmental, physiological, and activity sensors in a scalable, heterogeneous architecture. The SenQ query protocol provides real-time access to data and lightweight in-network processing. Circadian Activity Rhythm (CAR) analysis learns resident activity patterns and feeds them back into the network to aid context-aware power management and dynamic privacy policies.

## 1 Introduction

An aging baby-boom generation is stressing the U. S. healthcare system, causing hospitals and other medical caregivers to look for ways to reduce costs while maintaining quality of care. It is economically and socially beneficial to reduce the burden of disease treatment by enhancing prevention and early detection. This requires a long-term shift from a centralized, expert-driven, crisis-care model to one that permeates personal living spaces and involves informal caregivers, such as family, friends, and community.

Systems for enhancing medical diagnosis and information technology often focus on the clinical environment, and depend on the extensive infrastructure present in traditional healthcare settings. The expense of high-fidelity sensors limits the number that are available for outpatient deployment, and some require specialized training to operate. Manual record keeping has been identified as a key source of medical errors [1], and at its best, traditional data collection is intermittent, leaving gaps in the medical record.

Wireless Sensor Networks (WSNs) provide capabilities that are valuable for continuous, remote monitoring, as research into military and environmental systems attest. For healthcare applications, they can be deployed inexpensively

in existing structures without IT infrastructure. Data are collected automatically, enabling daily care and longitudinal medical monitoring and diagnosis. The wireless devices can integrate with a wide variety of environmental and medical sensors.

While addressing some of the needs of distributed healthcare, WSNs also present their own challenges to being practical, robust platforms for pervasive deployment. Privacy and security of collected medical data may be jeopardized by careless use of a wireless medium. Without smart power management, battery-powered sensors have short lifetimes of a few days or require continual maintenance.

We present AlarmNet, an Assisted-Living and Residential Monitoring Network for pervasive, adaptive healthcare in assisted-living communities with residents or patients with diverse needs. Contributions and novelties of this work include:

- An extensible, heterogeneous network middleware that addresses the challenges of an ad hoc wide-scale deployment, and integrates embedded devices, back-end systems, online analysis, and user interfaces;
- Novel context-aware protocols using two-way network information flow: environmental, system, and resident data flow into the back-end, and circadian activity rhythm analysis feeds back into the system to enable smart power management and dynamic alert-driven privacy tailored to an individual's activity patterns;
- SenQ, a query protocol for efficiently streaming online sensor data to the system and to users, integrated with privacy, power management, and activity analysis;

The system has been implemented and evaluated using custom and commodity sensors, an embedded gateway, and a back-end database and analysis program.

## 2 Related Work

An aging population has generated significant interest in smart environments that assist individuals. Several such systems and their differences with AlarmNet are described in this section.

University of Rochester's Smart Medical Home is a five-room lab outfitted with infrared sensors, computers, bio-sensors, and video cameras. The goal is to develop an integrated Personal Health System that collects data 24 hours a day and presents it to health professionals [2]. Similarly, Georgia Tech's Aware Home [3] is a prototype for an intelligent space that provides a living laboratory capable of knowing information about itself and its inhabitants' activities. It combines context-aware and ubiquitous sensing, computer vision-based monitoring, and acoustic tracking. AlarmNet relies less on installed infrastructure, allowing it to be deployed at larger scale and less expensively.

Massachusetts Institute of Technology (MIT) and TIAX, LLC are working on the PlaceLab initiative [4], which is a part of the House<sub>n</sub> project. PlaceLab is a one-bedroom condominium with hundreds of sensors installed in nearly

every part of the home. Researchers at Harvard have developed wireless sensors and software called CodeBlue [5] for medical applications, including pre-hospital and in-hospital emergency care and disaster response. The sensors include portable 2-lead ECG, pulse oximeter, wearable Pluto mote with built-in accelerometer, and a module with accelerometer, gyroscope, and electromyogram sensor for stroke patient monitoring. Researchers at University of Aarhus developed and deployed technologies for hospitals, including a location tracking system and two context-aware applications running on interactive wall displays and mobile phones [7].

AlarmNet combines the ubiquitous sensing of PlaceLab with the medical focus of CodeBlue and Aarhus, but in an assisted-living context where long-term (not just emergent) behavior and pathologies are learned. Also, devices in assisted-living and residential environments must be less obtrusive than in highly technological, expert-oriented hospitals.

University of Washington's Assisted Cognition project combines artificial intelligence and ubiquitous computing with clinical research on patient care. Their system uses relational Markov networks to model the complex relations among GPS readings, activities, and significant places [6]. As their focus has been on movement among places in a user's community, the work is largely orthogonal to AlarmNet, which targets the assisted-living residential population, and is primarily indoors where GPS is not reliable.

Another key difference from the state of the art is AlarmNet's two-way flow and feedback of information from back-end analysis processes to enhance system efficiency and enable autonomous, unobtrusive operation. One example is the Circadian Activity Rhythm (CAR) analysis program that learns the patterns of daily life of the residents. Dynamic privacy rules change when an individual exhibits a behavior that is critical to his health and enable authorized medical personnel to access vital data, which is otherwise restricted. Additionally, CAR enables advanced power management by anticipating which sensors should be kept active and which can be temporarily disabled to conserve power. In the future we can enhance other behavior-sensitive functions and services in the system with feedback and control from the CAR.

Another constructive aspect of the system is a flexibility that allows easy integration of new sensors or even mobile networks of sensors, and ad hoc deployment into existing structures. AlarmNet is a multi-platform, security- and privacy-aware architecture integrating a wide spectrum of heterogeneous devices such as wireless sensor nodes, gateways, PDAs, and back-end systems for data storage and analysis.

### **3 AlarmNet Architecture**

A key requirement for healthcare systems is the ability to operate continuously over long time periods and still integrate new technologies as they become available. AlarmNet satisfies these objectives by unifying and accommodating heterogeneous devices in a common architecture (see Figure 1) that spans wearable body networks, emplaced wireless

sensors, user interfaces, and back-end processing elements.

**Mobile Body Networks** are wireless sensor devices worn by a resident which provide activity classification or physiological sensing, such as an ECG, pulse oximeter, or accelerometers. The network is tailored to the patient's own medical needs, and can provide notifications (for example, alerts to take medicine) using an in-network wearable interface, the SeeMote we developed [8] with a color LCD.

We are integrating SATIRE [9], a body network that classifies Activities of Daily Living (ADLs) [10] by analyzing accelerometer data generated by a wearer's movements.

Body networks contain a designated gateway device that mediates interaction with the surrounding WSN. This modularizes the system's interaction with the body network to ease its integration. Data are streamed directly or multi-hop through the emplaced network to the AlarmGate gateways for storage, analysis, or distribution to user interfaces.

**Emplaced Sensors** are deployed in living spaces to sense environmental quality, such as temperature, dust, and light, or resident activities. Motion and tripwire sensors, in particular, provide a spatial context for activities and enable location tracking.

Due to their low-cost, small form factor, and limited power budget, the devices answer queries for local data and perform limited processing and caching. Though some deployment environments may enable the use of mains power, we do not require it so as to support ad hoc retro-fitting of existing structures. Figure 2(a) shows the lightweight stack resident on sensor devices.

AlarmNet supports dynamically adding new devices to the network, which register their capabilities and are initialized. This flexibility allows the system to change over time as sensors are developed or new pathologies require monitoring.

**AlarmGate** applications run on an embedded platform, such as the Crossbow stargate, and serve as a communication backbone and application-level gateway between the wireless sensor and IP networks. Owing to their greater resources, these devices perform major aspects system operation related to dynamic privacy, power management, query management, and security. The AlarmGate software stack is shown in Figure 2(b).

**Back-end** programs perform online analysis of sensor data, feeding back behavior profiles to aid context-aware power management and privacy. A database provides long-term storage of system configuration, user information, privacy policies, and audit records.

One such program, for Circadian Activity Rhythm (CAR) analysis, processes sensor data stored in the database and learns behavior patterns of residents. These are used to detect deviations from personal norms that may signal a short- or long-term decline in resident health.

The back-end is extensible to new analyses using a modular framework, wherein programs consume input sensor streams, filter and process them, and produce output streams in the database for other modules to use. These are

composed hierarchically from low-level sensor streams to high-level inference of symptoms and diseases.

**User Interfaces** allow doctors, nurses, residents, family, and others to query sensor data, subject to enforced privacy policies. We developed a patient-tracking GUI for a nurse’s station, and a query issuer for a PDA that graphs sensor data in real-time. These programs are not trusted components—they must connect through AlarmGate and do not have direct access to the database. This makes it easier to develop and deploy new interfaces customized to the application’s needs.

In summary, AlarmNet’s architecture supports health monitoring applications due to its flexibility and extensibility in 1) supporting dynamic addition of heterogeneous devices, sensors, and body networks, 2) feeding learned resident and system context back into the network, and 3) providing an open client model for future extension.

## 4 Query Management

A primary reason for developing AlarmNet was to use environmental, physiological, and activity data of assisted-living residents to improve their health outcomes. The automated analysis programs need to automatically collect data in the background, but the system must also support ad hoc queries by healthcare providers and the addition of new analysis programs over time. Existing data management solutions were mostly optimized for tree-based aggregation [11], or else used general-purpose virtual machines for arbitrary computation.

We developed *SenQ*, a query system that satisfies the requirements of our application domain: reconfigurable in-network sensing and processing, dynamic query origination by embedded devices, and high-level abstractions for expressing queries. A detailed treatment of SenQ’s lower layers and their performance was given in [12], with focus on sensing and in-network query processing. Here we present the system-level query management functions and their integration with other components of AlarmNet.

The back-end system, user interfaces, and embedded devices all issue queries using a common network protocol, in which queries are uniquely identified by  $\langle source\ ID, query\ ID \rangle$  tuples. Originators may request a snapshot of the current value or a periodic stream of a sensing modality. To reduce repetitive query parsing overhead on resource-constrained motes, both types of queries may be cached and efficiently restarted (or reissued) later.

Since radio communication in the WSN is expensive, we wish to process data at its source, if possible, to reduce the amount that is reported. However, sensor devices have limited memories and processors, and so only relatively lightweight processing is practical. SenQ dynamically constructs a scalar processing chain on the mote to perform spatial and temporal aggregation and filtering to reduce the energy consumed by communication.

**Query Manager** is a major actor in the query subsystem, and resides on the gateway, in the AlarmGate software. Devices are commonly added to and removed from the system, particularly in the health-care domain where monitoring needs evolve over time. To enable the Query Manager to maintain device state, nodes *register* with the nearest gateway

upon power-up, providing their device type, sensors, and hardware ID. They are assigned dynamic network IDs and localized via application-specific means.

The Query Manager issues background queries to devices as they are added to the network to satisfy the system's core management and tracking functionalities. Examples of background queries in AlarmNet are:

- All devices sample and report their battery supply voltage every four hours, but only if it is below 2.8V (indicating imminent failure);
- Motion, tripwire, and contact-switch sensors report activations on-demand, but no more often than every 100ms to debounce or dampen spurious bursts;
- Pulse oximetry devices, which are intermittently switched-on, collect heart rate and SpO<sub>2</sub> samples every 250ms, but report them every 750ms, each an average of three samples until the device is switched off; and,
- ECG sensors immediately begin reporting a stream of raw samples every 20ms, using full buffering to reduce network load and energy usage.

The Query Manager is the main point of access for *user interfaces*, translating between higher-level query abstractions and the SenQ protocol exchanged with particular sensor devices. Connected users receive a list of active devices that is updated in real-time as registrations are received. However, most users of the system will not have detailed knowledge of its current topology. Usability is improved if they can request sensor data semantically for people and locations. This presents a few challenges for query management on the gateway.

A request for sensor information about person  $P$  must be mapped to a device (or group of devices)  $D$  for execution. Some have static associations, such as a wearable-device owned or assigned to a user. Likewise for locations  $L$  in which fixed sensor nodes are placed. But since networks for assisted-living are more human-oriented and heterogeneous than most other types of WSNs, many sensor types require dynamic binding based on a person's context (location, activity, etc). Externalizing such bindings, as in other common approaches for WSNs, results in duplicated effort to track user and device state by both the query system and the application that uses it.

*Dynamic semantic binding* simplifies data access for users, but the challenge is how to provide it in a modular way that does not limit SenQ to a particular deployment environment. Our approach is to share the core of a context model, shown in Figure 3, with applications co-resident on the Query Management gateway. Relations among the core components (Device, Sensor, Location, Person, and Data) are maintained by SenQ and consulted for mapping queries to devices.

Instantiations of the system extend the model as appropriate; for example, AlarmNet adds privacy, power management, and activity analysis. These and any future extensions can access and mutate the core shared context to change

SenQ's bindings of queries to sensors. AlarmNet's context-aware modules are described next, beginning with activity analysis which also interacts with the others.

## 5 Circadian Activity Rhythms

It is known that most people exhibit behavioral trends in the home with 24-hour cycles, which can be called a "Circadian Activity Rhythm" (CAR). We have developed a CAR analysis program that measures the rhythmic behavioral activity of residents and detects changes within these patterns.

The CAR algorithm is statistical and predictive, and was first presented in [13]. Here we present its integration with AlarmNet and some recent extensions. We employ CAR in novel ways to improve both medical care and network performance. In particular, CAR supports context-aware protocols based on learned activity patterns for smart heterogeneous power management and dynamic alarm-driven privacy.

Rhythms based on an hourly distribution of the probability of user presence in every room are called *presence based CAR*. Those based on the density of the number of events per hour are called *activity level based CAR*. The CAR analysis program runs on a PC on the back-end of the system, and reads a database of resident activity.

The CAR program provides a GUI to display various information related to the activity analysis, such as the number of abnormal time periods (under-presence or over-presence in a room for presence based CAR), the degree of activity that occurred per hour and day during day or night (hypo or hyper-activity level for activity level based CAR), and the length and dates of stay of the resident.

Other graphs of the GUI display the main results of the CAR analysis. The graphs in Figure 4 present data from a clinical case study for a healthy resident who stayed 25 days in an assisted-living facility. The first one (Figure 4(a)) displays the average time the user spends in every room each hour, calculated over the number of days of the stay of the resident. On the right side, the graph indicates deviations in room presence. The experiments demonstrated that the CAR program needed a period of approximately two weeks to learn normal behavior patterns. The graphs in Figure 4(b) represent the same experiment (same subject, same clinical protocol, and same period of study) but for his activity levels.

Comparing both graphs gives complementary information concerning normal and abnormal activity levels in the different rooms of the dwelling. These graphs can provide a wealth of information about activity patterns such as the sleep/wake cycle, or some medical hints to the physician about some activities of daily living (ADLs) [10] of the resident such as eating, hygiene, and sleeping. In the future, more specific ADLs will also be inferred.

After the learning period, any statistically significant deviations from learned patterns are displayed as alerts in the GUI and are sent to the AlarmGate application. Nurses or physicians can investigate the source of the trouble by focusing on the region of the anomaly as identified by CAR. The hypothesis is that deviations from a resident's own

in-home behaviors can support medical diagnosis. For example, increased sleep time and multiple missed meals may indicate illness in the short-term, or reveal overall decline in the health of the resident in the long-term.

This hypothesis was investigated clinically in collaboration with the Medical Automation Research Center (MARC) at the University of Virginia medical school. Clinical behavioral patterns of older adults in assisted-living facilities were extracted from real data sets, and behavioral changes were studied by consulting the medical notebooks of the caregivers in charge of the monitored residents. The capability to detect anomalies from the norm calculated over multiple days was proved [14], but inferring pathologies or onsets of chronic pathologies remains to be explored.

## 6 Context-Aware Power Management

AlarmNet supports heterogeneous power sources in which some nodes are plugged into the wall and others operate on batteries. Stargates are plugged-in due to their high power consumption, while motes can be powered either way. This flexibility allows new battery-operated motes to be added quickly, in high density if needed, moved easily, and even placed outdoors where power is more problematic. Further, motes that are part of body networks necessarily use batteries (or scavenge energy from motion). Consequently, energy efficiency is an important design issue in AlarmNet.

Application demands of AlarmNet pose some particular requirements on power management. First, sensors are used to detect and collect information on residents, so they should adapt their operational states according to changes in the resident's behavior. Second, power management should provide openness to system administrators, who should be able to set policies unique for particular applications. Third, individual sensing modalities should be controllable, as well as radio components. For example, the system may want to set a high rate for temperature sensing, a low rate for light sensing, turn off other sensor types, and set a duty cycle mode for the radio. Last, in a heterogeneous network with diverse sensor nodes, such as ECG, motion, and weight sensors, power management should adaptively control each according to its own characteristics and context, including location.

To satisfy these requirements we designed a Context-Aware Power Management (CAPM) subsystem for AlarmNet. Many schemes in the literature save power using scheduling-based and/or sentry-based algorithms, often to maintain complete sensing coverage for unexpected events [15]. However, these schemes typically are not adjusted automatically or by administrators at runtime. Since CAPM is aware of the individualized patterns of residents, it dynamically chooses power management operations to provide good sensing service with less energy.

For battery-powered devices, our initial design of CAPM provides two types of power management operations: those based on administrator directives and those which rely on context-awareness.

First, administrators can directly control each sensor available on a mote. Sensors can be turned on/off or their rates set for each sensor type, and the radio is similarly controlled. Also, the effective period of each command can be set. A typical command may be, "Mote 1 turns off the light sensor, but senses the temperature every 2 seconds for the



following 2 hours.”

Second, CAPM tailors power management operations to the behavior patterns of the resident. Initially, administrators define some context policies for power management, such as “when the resident is sleeping or outside the apartment, turn off all sensors in the living room and reduce the temperature sensing rate to hourly in all rooms.”

Based on the resident’s behavior pattern learned by the CAR, the CAPM triggers the corresponding power management operation when the condition of a context policy is met. This context-aware operation provides more efficient power management that adapts to the resident’s behavior. Another advantage of context-aware operation is that the power subsystem is open for administrators to define their own context policies, according to application demands.

To resolve conflicting power management operations, we attribute to each a priority. Priorities are used to rank the importance of different commands and to interact with SenQ, which may request data from sensors that have been disabled. Background queries in AlarmNet have low priorities, while user-initiated streaming queries have high priorities.

The CAPM subsystem uses five functional components, shown in Figure 5: Sensor Drivers, a Context-Aware Power Manager, a Circadian Activity Rhythm (CAR) analysis program, a Context Manager, and User Interfaces.

**Sensor Drivers** control the sensing modalities available in motes, enabling, disabling or duty-cycling them in response to incoming power management directives. To manage conflicting commands, the sensor controller keeps a Sensor State Table (SST) that records the current state (on/off, sensing rate) and other information about the command which set the current state (such as its source and priority). Higher priority commands supercede lower priority ones, which are discarded without affecting the SST.

**Context Manager** collects and maintains the current context of residents and the environment from sensor data and back-end analysis. Context is the result of domain expert analysis based on sensor readings which indicate a resident’s health and environmental conditions.

**Context-Aware Power Manager**, residing in the AlarmGate application, controls power management operations for the system. This component is the core of CAPM. The inputs of the component are power policies defined by administrators, sensing reports from motes, the resident’s behavior patterns obtained from the CAR, and current system context. The output is a set of power management commands for each sensor device in AlarmNet.

We give a concrete example. The administrator may specify a context policy, “turn off all motion sensors when the resident goes to sleep,” and the CAR gives the resident’s activity pattern, “the resident usually sleeps from 10 P.M. to 6 A.M.” When context indicates that the resident is in the bedroom after 10 P.M., the context manager infers that the resident is sleeping, and sends power management commands to lower the duty-cycle of the sensors.

## 7 Dynamic Context-Aware Privacy

Data collected in AlarmNet reveals intimate details about a person's life activities and health status. As WSNs grow stronger in their capability to collect, process, and store data, personal information privacy becomes a rising concern. Our system includes a framework to protect privacy and still support timely assistance to residents in critical health situations.

Emergency-aware applications demand a privacy protection framework capable of responding adaptively to each resident's health condition and privacy requirements in real-time. Traditional role-based access control, which makes access authorization decisions based on users' static roles and policies, is not flexible enough to meet this demand. We designed a privacy protection framework which is dynamically adjustable to residents' context, allows data access authorization to be evaluated at runtime, and is able to adapt to residents' health emergencies.

A key novelty is that access rulings are dynamically altered based on context inferred by the CAR and other back-end programs, when necessary. For example, if a resident has blocked access to his ECG data for nurses, but the CAR has detected serious anomalous behavior that might indicate a heart problem, then a nurse is alerted and access to the data is allowed for a period of time.

In case of an alarming health status, privacy may be restricted or relaxed depending on context and who requests the data. This decision is recorded in the database with details of the mitigating context. It can be used to notify residents later about what transpired and why, and who accessed their data during the emergency.

Other context in our system includes the resident's physiological condition (ECG, pulse, blood pressure), living environment conditions (room temperature, light levels, dust), activities, and autonomy (inferred from ADLs by the CAR). System designers may specify privacy policies at different levels of granularity, from individual sensors to residents or groups of residents.

The privacy management framework resides in the AlarmGate application, and has three main functional components: the Context Manager, the Request Authorizer, and the Auditor (shown in Figure 6).

**Context Manager**, as described before, maintains residents' current locations, activities, and health conditions. In the privacy module, context is indexed by the tuple  $\langle \text{context id}, \text{context subject}, \text{context value} \rangle$ .

**Request Authorizer** is consulted when data queries are received at the Query Manager. It makes access decisions by consulting the system's privacy policies and the context of the query subject. After each access request is decided, it is recorded by the Auditor module.

SenQ allows queries for locations and devices, which must be mapped to a resident for authorization by the Request Authorizer. Locations may be assigned an owner, such as the resident living in a particular unit. Common areas without an owner use a default subject or policy. Devices also may have an owner, for example, an ECG that is worn by a particular resident. Otherwise, the device's assigned or current location is used to determine the context

subject.

**Auditor** maintains a trace of access requests in an audit trail, including the authorization decision made for each request (granted or denied). Dynamic decisions during alarming situations trigger the storage of details of the context for later notification of residents and investigation of misuse. Such records may be crucial in heavily regulated environments like health-care.

## 8 Testbed Implementation

We constructed a testbed to emulate the conditions of an assisted-living facility for ongoing evaluation of the overall system and of individual components.

AlarmNet uses a heterogeneous hardware environment of emplaced sensors, mobile body networks, gateways, and user devices. Figure 7 shows a sampling of the sensor and display hardware we implemented and integrated into AlarmNet using the lightweight embedded SenQ stack.

We use MicaZ and Telos Sky motes and X10 devices for the sensor platform. They are connected to a Crossbow Stargate SBC executing the AlarmGate application. Users access AlarmNet from PCs or iPAQ PDAs over the Internet or wirelessly from LCD-enabled motes.

Mobile body networks measure heart rate, pulse oximetry, heart electrical activity (ECG), and body movement. ECG and pulse oximeter sensors designed at Harvard for the CodeBlue project [5] have been integrated. SATIRE [9] uses five two-axis accelerometers worn by the resident to detect and classify activities. Other sensing modalities include body weight, systolic and diastolic blood pressure, motion, dust, light, optical tripwires, and magnetic reed switches. These provide rich inputs to the system for monitoring residents' physiological and environmental conditions.

Software modules shown in Figure 2(a) are implemented in nesC for TinyOS 1.1.15 on the motes, and in Java for the Blackdown v1.3.1 JRE for ARM-Linux. Back-end programs include MySQL 5.0, the CAR program in NI LabWindows/CVI, and others in Python.

User interfaces have been implemented in Java, including a resident tracking GUI and administration interface (see Figure 8, right side), suitable for a nurse's station. It shows a floor-plan of the living space and indicates the current location of the resident and sensor devices.

A Query Issuer is implemented in Java and runs on the IBM WebSphere Micro Environment JVM, on an HP iPAQ 5550. It allows a user to connect to the AlarmGate, authenticate as a valid user, and issue queries. Sensor data are graphed in real-time as reports stream in from the network.

Individual components of the system have been evaluated [8, 9, 12, 14], and full-system evaluation is ongoing.

## 9 Conclusion

AlarmNet is a wireless sensor network system designed for long-term health monitoring in assisted-living environments with two central design goals. First, system operation adapts to the individual context and behavior patterns of the residents, which feeds back to influence power management and privacy policy enforcement. Second, the system is extensible and supports a diverse collection of sensors, user interfaces, and power and privacy policies.

## References

- [1] President's Information Technology Advisory Committee (PITAC), Revolutionizing health care through information technology, June 2004.
- [2] A. Almudevar, A. Leibovici, and C. Horwitz. Electronic motion monitoring in the assessment of non-cognitive symptoms of dementia. In *12th International Congress of the International Psychogeriatric Association*, 2005.
- [3] C. D. Kidd, R. J. Orr, G. D. Abowd, C. G. Atkeson, I. A. Essa, B. MacIntyre, E. Mynatt, T. E. Starner, and W. Newstetter. The aware home: A living laboratory for ubiquitous computing research. In *Proceedings of the Second International Workshop on Cooperative Buildings*, 1999.
- [4] S. S. Intille, K. Larson, E. Munguia Tapia, J. Beaudin, P. Kaushik, J. Nawyn, and R. Rockinson. Using a live-in laboratory for ubiquitous computing research. In *Proc. of PERVASIVE*, pages 349–365, 2006.
- [5] D. Malan, T. Fulford-Jones, M. Welsh, and S. Moulton. Codeblue: An ad hoc sensor network infrastructure for emergency medical care. In *Proceeding of the International Workshop on Wearable and Implantable Body Sensor Networks*, 2004.
- [6] L. Liao, D. Fox, and H. Kautz. Location-based activity recognition using relational markov networks. In *Proc. of IJCAI*, 2005.
- [7] J. E. Bardram, T. R. Hansen, M. Mogensen, and M. Soegaard. Experiences from real-world deployment of context-aware technologies in a hospital environment. In *Proc. of Ubicomp*, pages 369–386, 2006.
- [8] L. Selavo, G. Zhou, and J. A. Stankovic. SeeMote: In-situ visualization and logging device for wireless sensor networks. In *Proc. of BASENETS*, 2006.
- [9] R. K. Ganti, P. Jayachandran, T. F. Abdelzaher, and J. A. Stankovic. SATIRE: A software architecture for Smart AtTIRE. In *Proc. Mobisys*, 2006.
- [10] S. Katz, A. B. Ford, R. W. Moskowitz, B. A. Jackson, and M. W. Jaffe. Studies of illness in the aged. the index of ADL: A standardized measure of biological and psychosocial function. *JAMA*, 185:914–919, 1963.
- [11] S. Madden, M. Franklin, J. Hellerstein, and W. Hong. TinyDB: An Acquisitional Query Processing System for Sensor Networks. *ACM TODS*, 30(1):122–173, 2005.
- [12] A. D. Wood, L. Selavo, and J. A. Stankovic. SenQ: An Embedded Query System for Streaming Data in Heterogeneous Interactive Wireless Sensor Networks. In *Proc. DCOSS*, 2008.
- [13] G. Virone, N. Noury, and J. Demongeot. A system for automatic measurement of circadian activity deviation in telemedicine. *IEEE Transactions on Biomedical Engineering*, 49(12):1463–1469, December 2002.

- [14] G. Virone, M. Alwan, S. Dalal, S. Kell, J. A. Stankovic, and R. Felder. Behavioral patterns of older adults in assisted living. To appear in *IEEE Transactions on Information Technology in Biomedicine*.
- [15] D. Tian and N. D. Georganas. A node scheduling scheme for energy conservation in large wireless sensor networks. *WCMC*, 3(1):271–290, 2003.

## Biographies

Anthony D. Wood is a Ph.D. candidate in the Department of Computer Science at the University of Virginia, where he recently won the Outstanding Graduate Research Assistant Award. His research interests are wireless sensor networks, security, networking, and embedded computing, particularly as applied to healthcare applications. Anthony received an M.C.S. from the University of Virginia and a B.S. from Virginia Tech.

Professor John A. Stankovic is the BP America Professor in the Computer Science Department at the University of Virginia. He is a Fellow of both the IEEE and the ACM. He also won the IEEE Real-Time Systems Technical Committee's Award for Outstanding Technical Contributions and Leadership (inaugural winner). He also won the IEEE Technical Committee on Distributed Processing's Distinguished Achievement Award (inaugural winner). His research interests are in distributed computing, real-time systems, operating systems, and wireless sensor networks. Prof. Stankovic received his Ph.D. from Brown University.

Gilles Virone...

Leo Selavo is an Associate Professor at the Department of Computer Science of the University of Latvia. His research interests are wireless sensor networks and embedded systems for healthcare, environmental sciences and mobile applications. Leo received a Ph.D. in Computer Science from the University of Pittsburgh in 2004. Subsequently he was a Research Associate at the University of Virginia where he participated in AlarmNet and Luster projects.

Zhimin He is a Software Engineer at Google Inc., where he develops the Google AdWord Serving System. His research interests are systems and data mining. Zhimin He received a Master in Computer Science from the University of Virginia in 2004.

Qihua Cao is a Ph.D. student in the Department of Computer Science at the University of Virginia. Her research interests are wireless sensor networks, self-healing, maintenance-in-field, distributed computing, and real time systems. She actively contributed to diverse projects, such as a military application VigilNet, an age-at-home application AlarmNet, and an environmental monitoring application Luster, in which she authored and co-authored several core papers.

Thao Doan is a Program Manager at the Microsoft Corporation. She received her Masters degree in Computer Science from the University of Virginia in 2006. Her research interests are information privacy and wireless sensor networks.

Yafeng Wu is a Ph.D. candidate at the University of Virginia. His research interests are in embedded systems, wireless sensor networks, and ad hoc networks. Yafeng received his Master and Bachelor degree in Computer Science from the University of Science and Technology of China.

Lei Fang...

Radu Stoleru is an Assistant Professor in the Department of Computer Science at Texas A&M University. His research interests are in deeply embedded wireless sensor systems, distributed systems, embedded computing and computer networking. He has authored over twenty-seven papers and won the Outstanding Graduate Student Research Award from the Department of Computer Science, University of Virginia in 2007. Radu Stoleru received a Ph.D. in Computer Science from the University of Virginia in 2007.

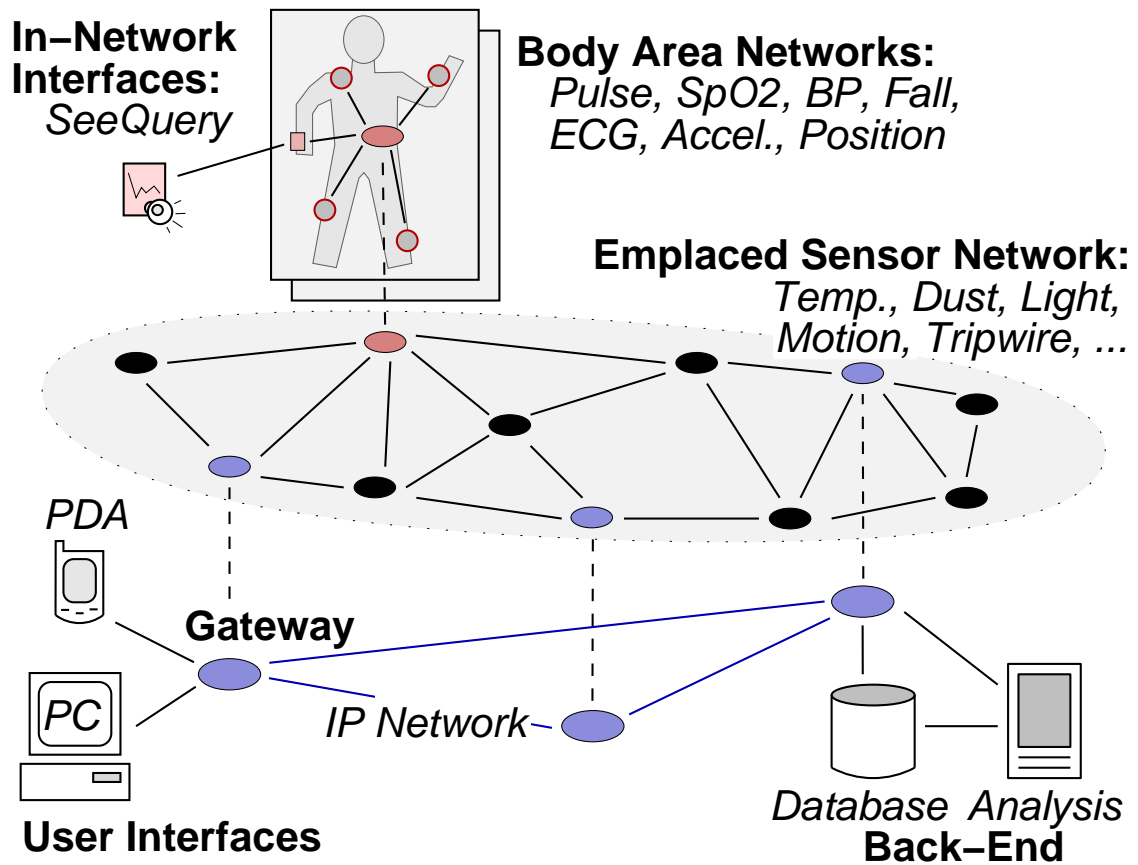


Figure 1: AlarmNet architecture.

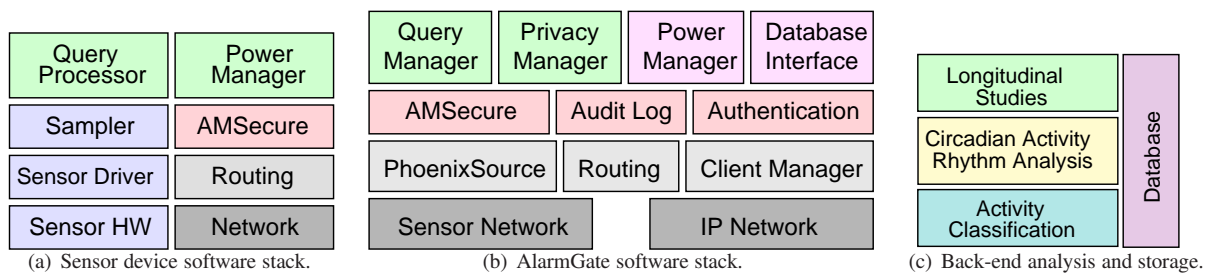


Figure 2: AlarmNet software architecture for sensor devices, AlarmGate gateways, and back-end servers.



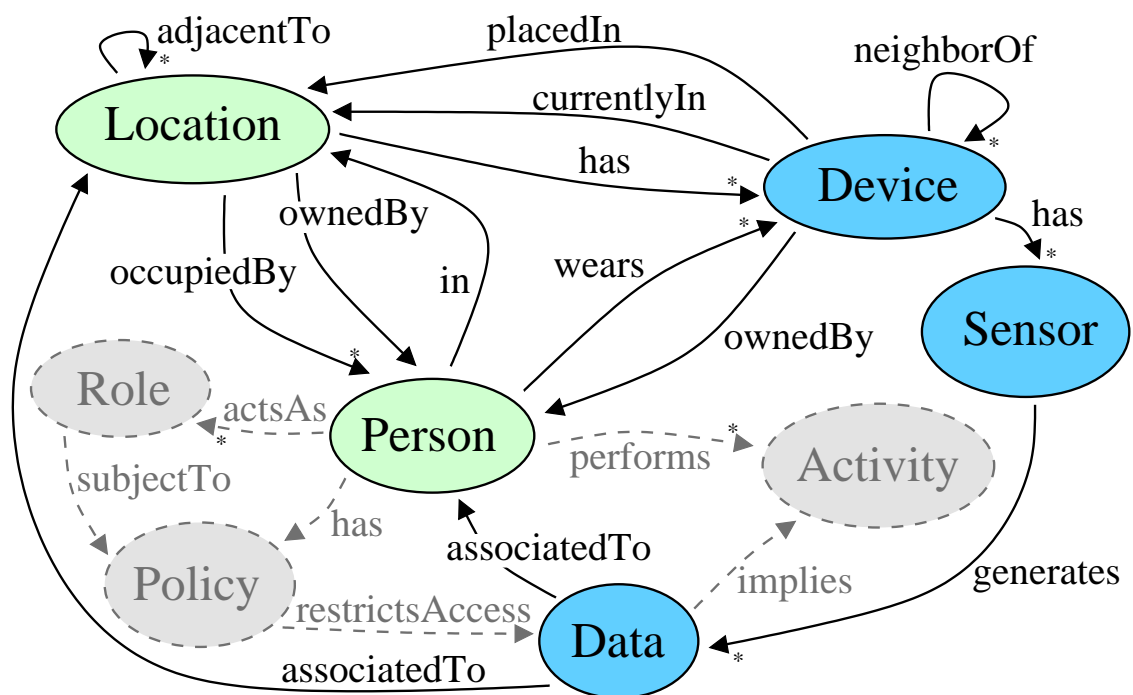



Figure 3: SenQ's context model. AlarmNet extensions are in grey.



(a) Presence based CAR analysis GUI. Average time spent in every room per hour is graphed on the left side.

(b) Activity level based CAR analysis GUI. Average number of motion sensor events in every room per hour is on the left side.

Figure 4: Circadian Activity Rhythm (CAR) analysis GUIs. Sums of daily deviations from the user's norm are on the right side of each GUI, showing a learning period after initial deployment.

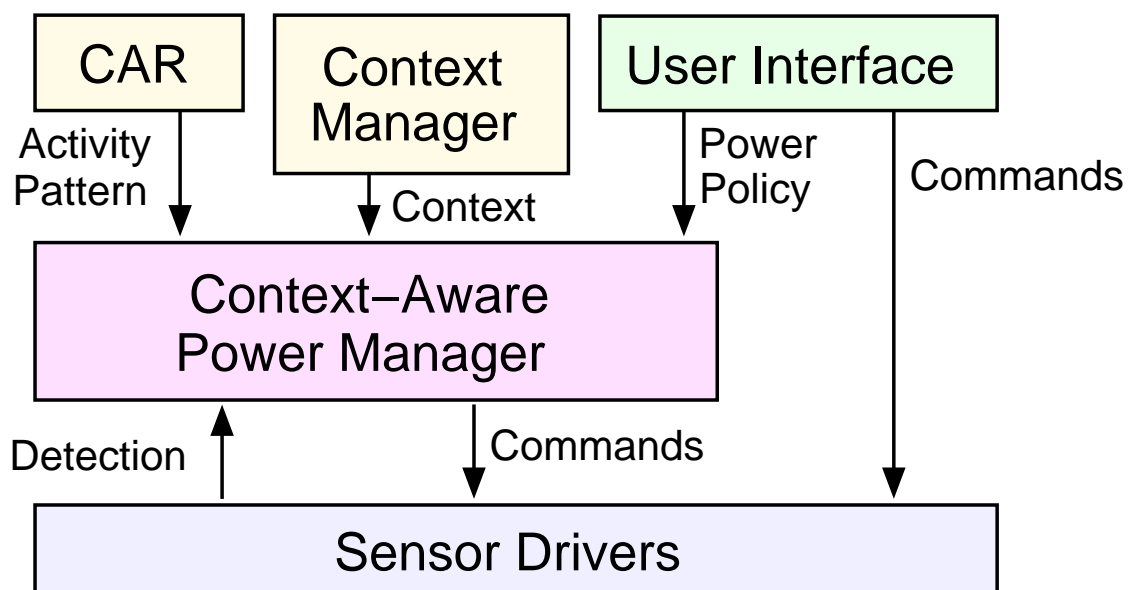


Figure 5: CAPM components in AlarmNet.

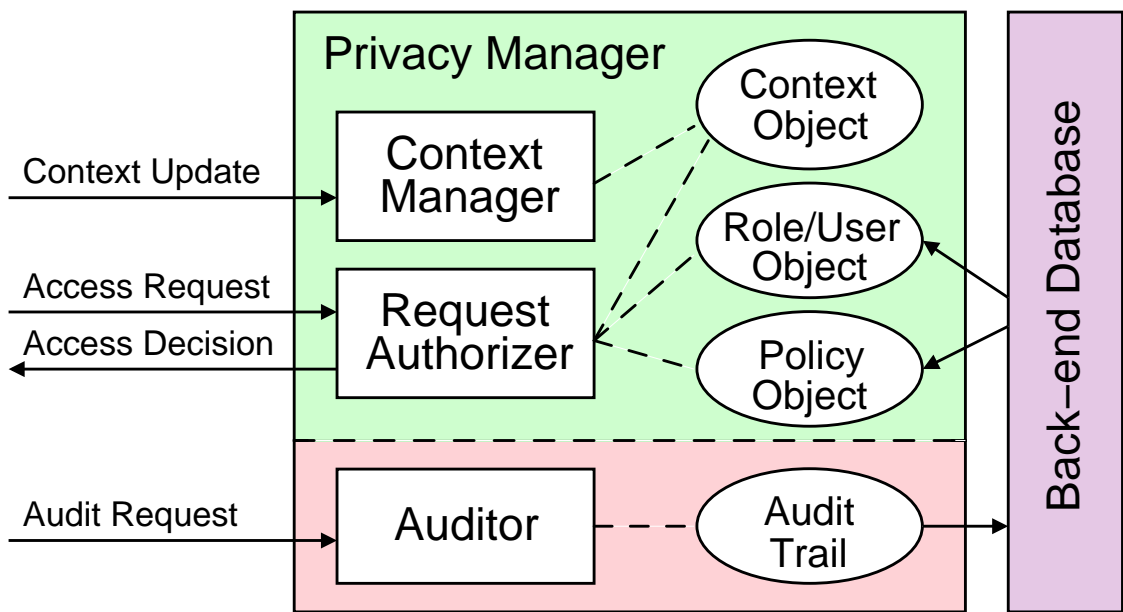


Figure 6: Privacy-related components in AlarmNet.



Figure 7: Sensor and display hardware in AlarmNet.

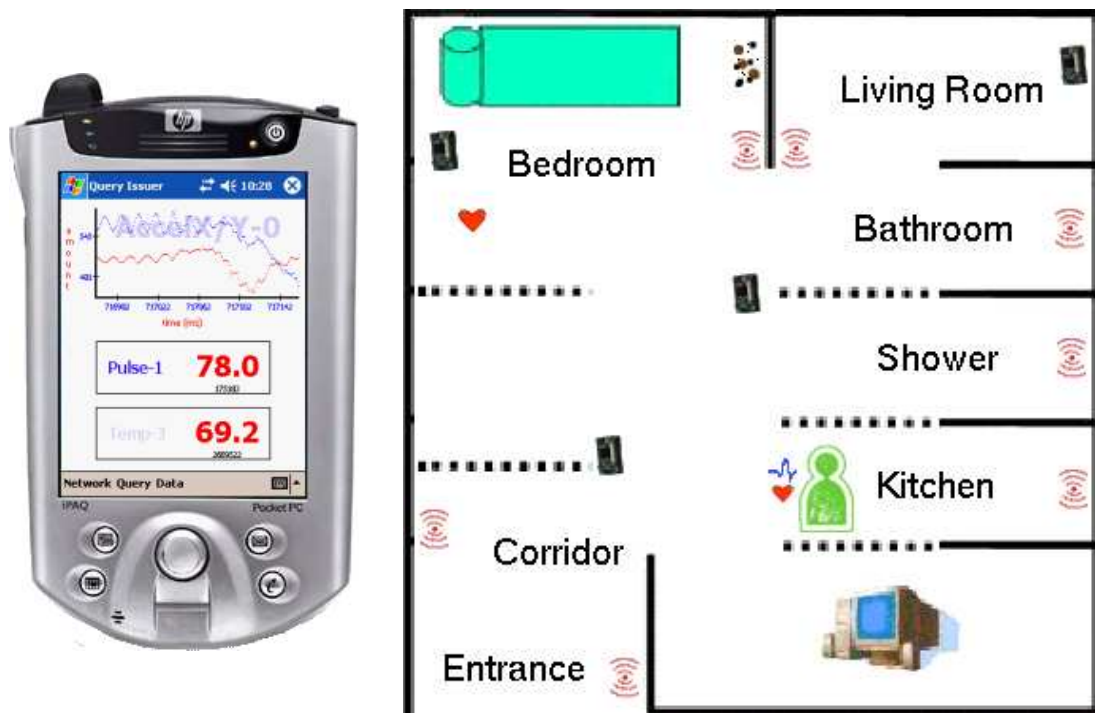


Figure 8: AlarmNet user interfaces: PDA query issuer and nurse's tracking and administration GUI.