Research Challenges in Computation, Communication, and Context Awareness for Ubiquitous Healthcare

Hariharasudhan Viswanathan, Baozhi Chen, and Dario Pompili, Rutgers University

ABSTRACT

A new paradigm for ubiquitous healthcare characterized by pervasive continuous vital sign data collection, real-time processing of monitored data to derive meaningful physiological parameters, and context-aware data- and patient-centric decision making, is central to deliver personalized healthcare solutions to the elderly and the physically challenged. However, this new paradigm requires real-time processing of wirelessly collected vital signs using inherently complex physiological models and analysis of the processed information under context (e.g., location, ambient conditions, current physical activity) to extract knowledge about the health condition of patients. As the computational capabilities of biomedical sensor nodes are insufficient to run these models, this article presents an innovative resource provisioning framework that organizes and harnesses the computing capabilities of under-utilized electronic devices in the vicinity (e.g., laptops, tablets, PDAs, DVRs, medical terminals) in home and hospital settings. Novel wireless communication solutions for reliable vital sign transmission and algorithms for acquiring context awareness to support this framework are also discussed.

INTRODUCTION

Developed countries currently face major problems in economically sustaining their healthcare programs due to the rising costs of hospitalization and specialized institutions. At the same time, rural and inaccessible areas of the *developing countries* suffer from a major problem in terms of high to very high patient-to-doctor ratios;[1] approximately 1,500:1 in Latin America and South Asia, 5,000:1 in parts of South-East Asia and Oceania, and 20,000:1 in Central Africa. These two seemingly disparate problems, however, highlight a common and immediate need for a new paradigm for ubiquitous healthcare characterized by

- Pervasive continuous vital sign data collection;
- Real-time seamless access and processing of monitored data to derive meaningful physiological parameters;
- Context-aware data- and patient-centric decision making.

The new paradigm is central to assist healthcare personnel in delivering personalized care for the elderly and the physically challenged in developed countries as well as in providing better medical attention and intervention for currently large under-served populations in developing countries.

The rapid growth of wireless and non-invasive sensor technologies and low-power wireless communication technologies [1] has enabled continuous monitoring of mobile patients using compact biomedical sensor nodes, also called *motes*. These small wearable devices — limited in memory, energy, computation, and communication capabilities — are capable of continuously monitoring vital signs such as blood pressure and flow, temperature, electrocardiogram (ECG), electromyogram (EMG), oxygen saturation, and CO_2 concentration. However, there are numerous tasks involved in rendering this data useful.

First, significant amount of preprocessing like noise rejection, data disambiguation, and consistency check has to be performed on the raw data provided by the failure-prone inexpensive sensors. Second, there is a need to prioritize the transmission of huge amount of preprocessed vital sign data from multiple body sites on a patient and from multiple patients (say, in a retirement home or hospital setting) in terms of their degree of importance for diagnosis to avoid network traffic congestion and to maximize reliability. Third, the collected vital signs have to be input into compute-intensive models to derive meaningful physiological parameters of interest. For example, heart rate variability (HRV) is a parameter that can be derived from preprocessed ECG data using frequency- or time-domain analysis. HRV is of greater interest to physicians than the raw ECG as it is a marker of biological aging as well as an indicator of increased risk of cardiovascular diseases and premature mortality.

The success of a ubiquitous healthcare system lies in how intelligently it can take into account the user's context, for example, how well is the acquired vital sign data correlated with the user's current situation. Context awareness is defined

[1] *Patients per doctor map: http://bigthink.com/ideas/21237*

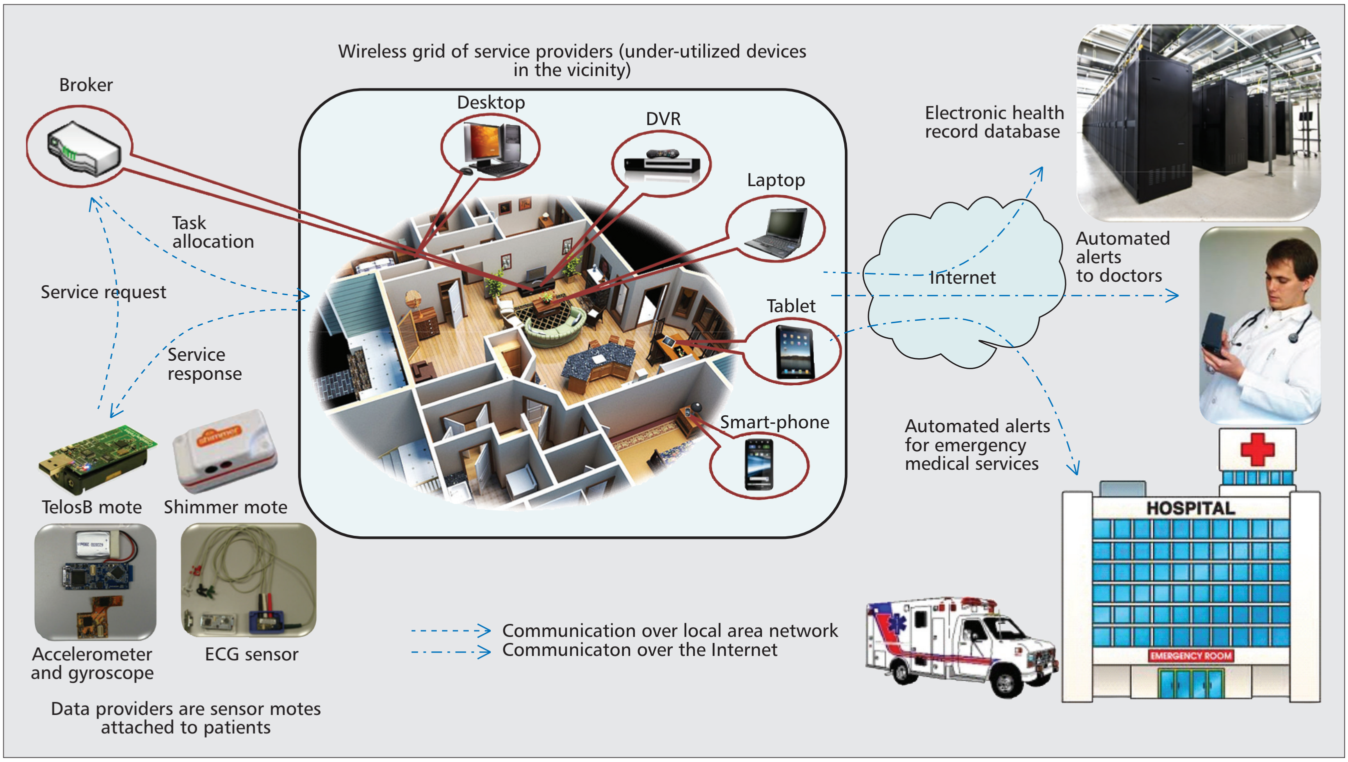

Figure 1. *Concept diagram illustrating our envisioned ubiquitous healthcare solution.*

as the state of knowledge of external and internal entities that causes a change in the user's situation, thus necessitating a different interpretation of the data in hand [2]. In the case of ubiquitous healthcare, contextual information such as physical activity, location, and environmental conditions can be obtained by processing data from kinematic sensors (accelerometers and gyroscopes), GPS, and temperature/humidity sensors, respectively. Correlation of derived physiological parameters with contextual information can be used for automated diagnosis of physiological conditions. For example, insights into the cardiac health of a subject can be acquired by correlating the HRV with contextual information such as state of exertion or rest identified using physical activity recognition, which employs real-time distributed sensing and compute-intensive machine-learning-based techniques. Such a system can also be used to measure the effectiveness of active physiotherapy, to perfect techniques of sport persons, and to monitor the elderly remotely and trigger emergency response.

It is clearly evident that simultaneously executing algorithms aimed at preprocessing and prioritizing raw patient data as well as at running models for deriving physiological parameters and for acquiring context awareness requires computing capabilities that go beyond those of wearable sensor motes. To overcome this limitation, we envision that the computing and storage capabilities of under-utilized electronic devices in the vicinity such as desktop and laptop computers, digital video recorders (DVRs), mobile and static medical terminals, tablets, and smart phones (at homes and hospitals) can be utilized to form an *elastic resource pool* that can process massive amounts of locally generated vital sign data in parallel.

The proposed approach will significantly improve the response time, quality, and relevance of data- and compute-intensive medical applications. Our emphasis on the proximity of data and computation is due to the prohibitive communication cost and response time involved in enabling these applications using the *cloud-computing* approach — in which computation and storage are offloaded to *remote servers on the Internet* called the *cloud*. Aggregating large amounts of sensor data at a remote location for centralized computation incurs huge communication overhead leading to significant delays [3]. This approach also imposes the unrealistic requirement of reliable as well as continuous connectivity to the cloud (via a gateway). In the reminder of the article, we discuss our envisioned solution for ubiquitous healthcare, the associated research challenges that have to be addressed, and our proposed innovations for the same.

Envisioned Approach, Challenges, and Contributions

Our envisioned solution for ubiquitous healthcare involves organizing powerful devices in the vicinity into a *wirelessly connected local mobile computing grid* as depicted in Fig. 1. The collective computational capability of this mobile grid can be exploited to perform in-network prioritization of vital sign data, to derive meaningful physiological parameters from vital signs using models, and, finally, to improve drastically the accuracy of early diagnosis by understanding the derived parameters under context.

Our proposed resource provisioning framework strives to minimize the computational load on individual service providers by exploiting parallelism while incurring minimal communication cost.

The sensor motes, acting as *data providers*, can offload the task of executing compute-intensive algorithms and models to these networked static and mobile devices, the *service providers* (depicted in Fig. 1). Vital sign data can be transmitted wirelessly from the data providers to service providers using Bluetooth or ZigBee. A *broker* is responsible for managing the elastic resource pool (i.e., handling service requests from data providers and allocating tasks to available service providers). Service requests are workloads (e.g., compute-intensive physiological models, algorithms for data disambiguation, models for acquiring context awareness) with varied computational and storage requirements. The broker can be a separate entity or a logical role played by one of the service providers. The information obtained by processing vital sign data in the local computing grid can be uploaded to an electronic health record database, and can also be used to generate and send alerts to emergency medical service personnel over the Internet. However, there are numerous research challenges associated with realizing our envisioned solution for ubiquitous healthcare.

First, discovery and provisioning of computing resources are major research challenges in wireless grid computing. This is because the device availability is highly unpredictable and device mobility is also very high. To address this challenge, we propose a *resource provisioning framework* for mobile grids that runs at the broker. The framework facilitates interactions between data providers, which place service requests, and service providers, which dedicate a portion of their computational (CPU cycles), storage (volatile and non-volatile memory), and communication resources (i.e., network interface capacity) for servicing those requests.

Second, prioritization and reliable transmission of patient vital signs to the wireless grid is non-trivial. Apart from the radio frequency (RF) interference produced by transmissions from multiple patients in homes and hospitals, proliferation of numerous electronic devices using the same radio technologies has resulted in additional RF interference. To address this challenge, we propose a quality of service (QoS)- and interference-aware wireless body area network (BAN) system [4] that includes an intra-BAN protocol, an inter-BAN protocol, and a packet scheduling scheme. While the intra-BAN protocol mitigates interference within a BAN, the inter-BAN protocol employs channel-quality-aware medium access and cooperative multi-hop routing to ensure connectivity with the local computing grid. The packet scheduling scheme prioritizes data streams based on the severity of the patient's condition to avoid congestion and maximize reliability. Reliability is the ratio of the number of data packets received by the base station (usually a medical terminal) and the number of packets sent by a source node.

Third, while obtaining the environmental conditions and physical location does not involve significant processing of sensor data, recognizing human physical activities for context awareness requires the design and development of novel distributed algorithms that process the data provided by the kinematic sensors attached to different body sites. We cannot adopt existing camera- and computer-vision-based techniques as they are not pervasive, may interfere with the day-to-day lifestyle of the user, and do not maintain user privacy. To address this challenge, we propose a novel pervasive *windowing- and learning-based technique*[2] [5] to recognize on the fly human physical activities, the specific movements comprising those activities, and their start and end times. All of our proposed innovations overcome the limitations in the state of the art and hence enable realization of our envisioned solution for ubiquitous healthcare in developing and developed countries.

Proposed Resource Provisioning Framework

Our proposed resource provisioning framework strives to minimize the computational load on individual service providers by exploiting parallelism while incurring minimal communication cost. The framework applies to medical applications exhibiting *data parallelism* (in which data is distributed across different parallel computing nodes that perform the same task) as well as to medical applications exhibiting *task parallelism* (in which parallel computing nodes may perform different tasks on the same or different data). Additionally, it can also facilitate storage of the collected vital signs and processed contextual information in remote electronic health record databases (as shown in Fig. 2) to provide crucial historical data and information whenever required for data- and patient-centric decision making.

In our framework, the broker is composed of two components, namely, *workflow manager* and *scheduler/optimizer* as shown in Fig. 2. The workflow manager (also called the *master*) handles service requests from data providers. The optimizer identifies the number of service providers (also called *workers*) available in the vicinity for the requested duration and allocates workload tasks among them. The broker is aware of the availability of workers either through voluntary service advertisements from the workers or through a request-and-response mechanism. The advertisements include information about the amount of computing (in terms of CPU cycles), memory (in bytes), and storage (in bytes) resources as well as the duration for which they are available for use by the data providers. Communication among the master and the workers as well as among the optimizer and workers happens over Comet,[3] a scalable peer-to-peer content-based coordination space developed by our colleagues at the NSF Cloud and Autonomic Computing Center, Rutgers University. The messages in this coordination space (also called an information space) are constructed in the form of tuples (XML strings) as shown in Fig. 2.

The broker shares the different tasks of the workload submitted by the data providers among the available service providers based on one of several possible policies. An example policy may be minimization of response time with emphasis on the proximity of data and computation while ensuring that none of the service providers is unfairly overloaded. Another policy might just

[2] *Activity Recognition Demo: http://nsfcac.rutgers.edu/CPS/projects/ban/*

[3] *Comet: http://nsfcac.rutgers.edu/CometCloud*

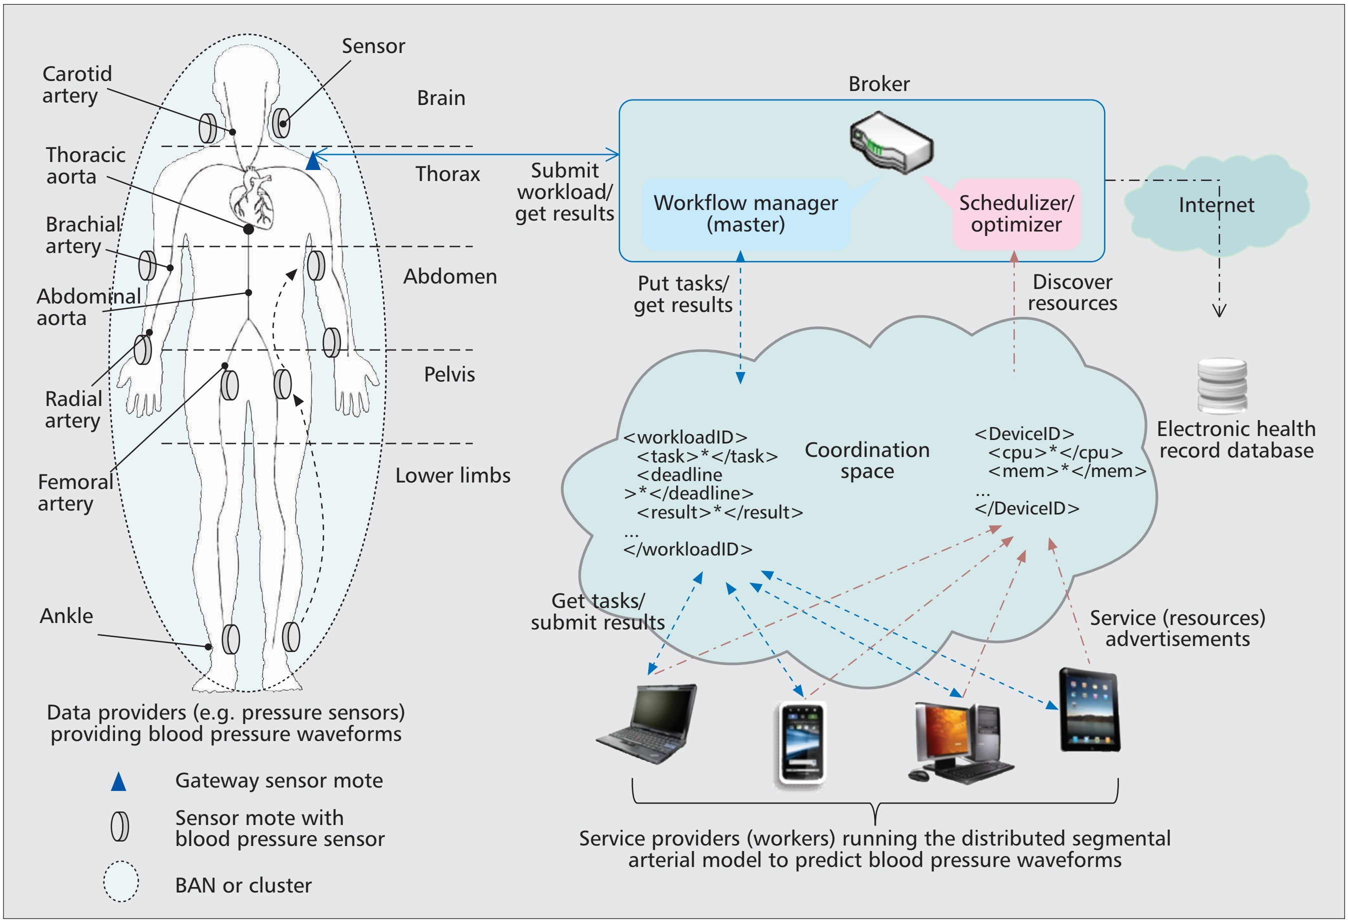

Figure 2. *Efficient internal bleeding compartment estimation using pressure sensors (data providers), computing devices (service providers), and the resource provisioning framework (on the broker).*

place emphasis on response time without considering fairness. While the under-utilized electronic devices at home can be configured or setup to be service providers by default (as suggested by [3]), mechanisms have to be formulated to incentivize mobile devices that are part of the network in retirement-homes and hospitals to play the role of service providers. Prior efforts on mobile grids for ubiquitous healthcare focused primarily on sensing or data collection. In [6] (mHealth), the authors propose mobile-phone-based architectures for social and behavioral sensing. These architectures make accurate monitoring of subjects feasible by enabling ubiquitous collection, storage, and sharing of relevant data. However, they overlook the key issue of real-time data processing for just-in-time treatment/intervention and context-awareness. In contrast, our proposed approach exploits mobile devices as data and service providers, handles uncertainty in the mobile environment to impart robustness, and ensures the QoS of real-time applications even under highly dynamic and unpredictable operating conditions. Additional results, details on the state of the art in mobile grid computing, and a demo of our Android-based prototype of the resource provisioning framework powering a distributed data-parallel compute-intensive application in tablets can be found in the Cyber Physical Systems (CPS) Laboratory webpage.[4]

To appreciate the effectiveness of using a local computing grid and our proposed framework for real-time processing of vital signs, consider the following example. Injuries leading to hemorrhagic shock (in particular, internal bleeding) if undetected and unattended to during the so-called golden hour [7], which is usually the first hour after the injury, may result in death. A reliable, accurate, and clinically useful solution for locating the site of hemorrhage is online processing and analysis (using physiological models) of vital sign data (blood pressure) collected simultaneously from multiple body sites. Our collaborators at the Department of Biomedical Engineering, Rutgers University, have developed computational models to predict pulsatile blood pressure waveforms in large arteries that can aid in early detection of internal bleeding, occlusions, arterial stenosis, and other altered vasoactive conditions [8]. However, frequency domain distributed segmental arterial models, which employs wave reflection theories, as well as other existing models that employ transmission line theories or non-linear differential equations, are all heavily compute intensive. As shown in Fig. 2, non-invasive blood pressure sensors attached to different body sites can serve as the data providers while the devices in the vicinity can serve as service providers. The collective computational capabilities of the service providers can then be exploited to execute the com-

[4] *Mobile grid demo: http://nsfcac.rutgers.edu/CPS/projects/fareshare*

The primary needs of an ubiquitous healthcare system are reliable wireless transmission of patient vital signs and prioritization of data streams within and across body area networks based on the severity of the patient's condition to avoid network traffic congestion and to maximize reliability.

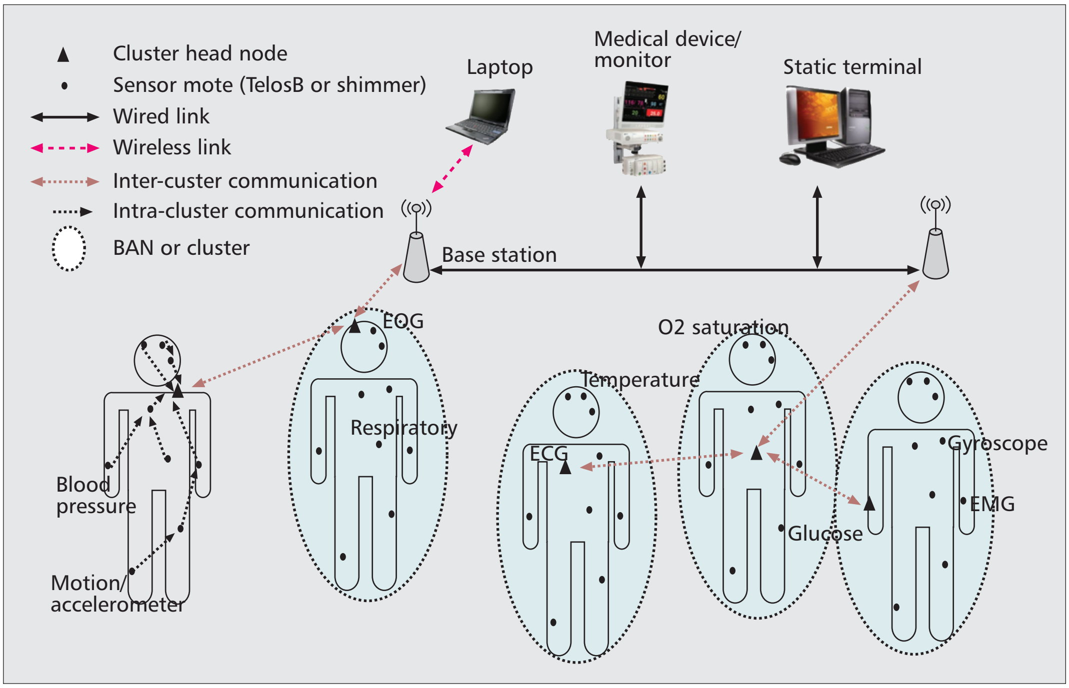

Figure 3. *Network architecture for reliable continuous vital sign monitoring.*

pute-intensive models and estimate the blood pressure waveforms that will provide detailed insights into the cardiovascular health of the patient.

Prioritization and Reliable Transmission of Vital Signs

The primary needs of an ubiquitous healthcare system are reliable wireless transmission of patient vital signs and prioritization of data streams within and across BANs based on the severity of the patient's condition to avoid network traffic congestion and to maximize reliability. Existing technology in home as well as in pre- and hospital environments lacks effective methods for prioritizing multiple data streams, for evaluating time-dependent trends, for managing incomplete data, and for providing effective alerts because they have mainly focused on the design of one-BAN systems (e.g., [9, 10]).

We propose a novel QoS-aware wireless communication solution that (with minimum added RF interference):

- Collects and prioritizes multiple vital sign transmission from different BANs in retirement home and hospital settings;
- Reliably transfers the acquired patient data to the service providers for processing;
- Seamlessly supports patient mobility by ensuring continuous connectivity between the BANs and the wireless grid through multihop routing via neighboring BANs.

Reliable transmission of vital signs is of paramount importance to ensure reliable performance of physiological models that require error-free data for processing and diagnosis. Also, prioritization of data streams originating from multiple sensor motes of one BAN and of different BANs based on the severity of the patient's condition (called *triage* [4]) will assist medical personnel when the resources are insufficient for all patients to be treated immediately.

The proposed networking system is designed to be a *two-tier hierarchical architecture* with the high-level tier composed of BANs (networked via inter-BAN communications) and the low-level tier composed of sensor nodes in each BAN (networked via intra-BAN communications) as shown in Fig. 3. Every BAN is formed by wirelessly networked sensor nodes that collect and transmit patient physiological data. Inside each BAN, a distributed cluster head (CH) selection mechanism is adopted in order to elect a node with higher computing, energy, and networking capabilities to play the logic role of a gateway node. This CH collects, aggregates, and fuses the data from other sensors in the cluster and transmits via multihops the fused data to the best base station using other CHs (in different BANs) as relay nodes. The wireless station will then relay the data to static or mobile service providers in the computing grid for pre-processing or as inputs to physiological models.

Inter-BAN communications: Our solution adopts a *cross-layer modular design*, which includes medium access control (MAC), packet routing, and packet scheduling functionalities. Wireless sensor motes have the capability to select one of the multiple available frequency channels to communicate with each other. In our solution, the link/channel quality indicator (LQI) is considered in the MAC and routing modules, which provides the foundation for our interference-aware multichannel quality-based MAC (MQMAC) and channel-quality-based routing (CBQR) protocols. The quality of all the channels is estimated using a probing mechanism and the channel with the best quality is determined for use via a two-way handshaking mechanism involving both the sender and the receiver. The

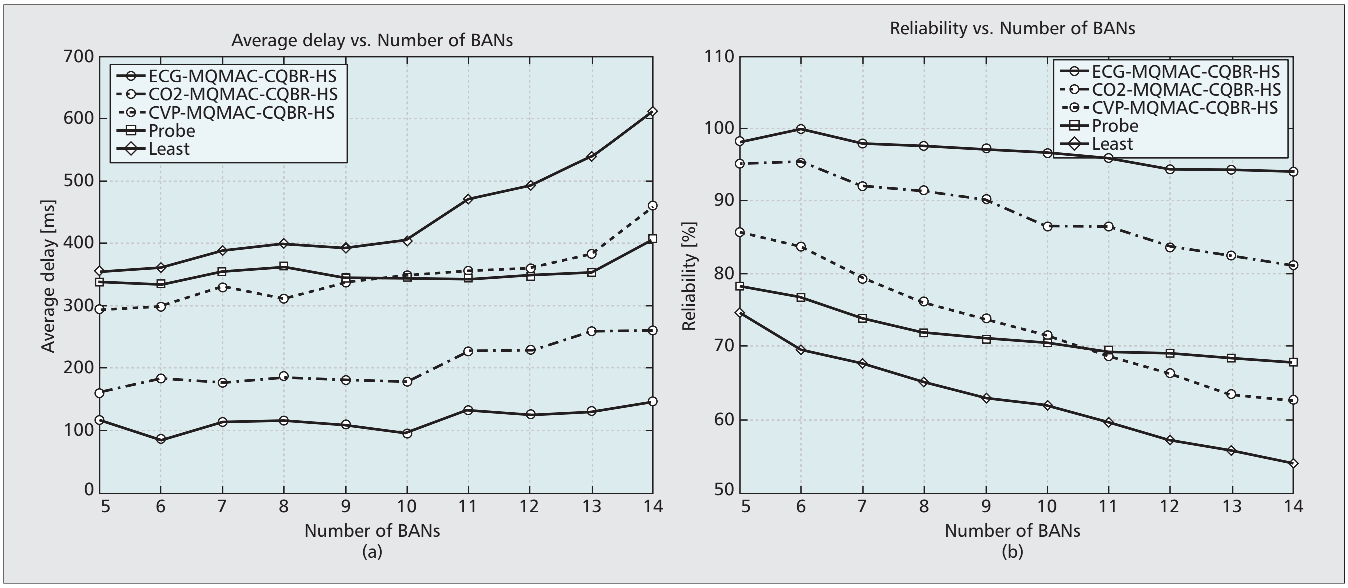

Figure 4. *Performance analysis in terms of average delay and reliability.*

channel quality is also considered in the routing module, i.e., paths with low quality links are not selected. Moreover, a two-level packet scheduling scheme is proposed to maximize the reliability for various traffic classes while guaranteeing their end-to-end (e2e) requirements.

Intra-BAN communications: The aim of the intra-BAN protocol is to facilitate cooperation among the sensor motes belonging to one BAN to perform jointly non-compute-intensive preprocessing of raw data (for extracting meaningful features, filtering, avoiding redundancy, etc.) while maintaining the amount of aggregated interference generated by the BAN at a bare minimum. As we already pointed out, compute-intensive preprocessing and execution of complex models is offloaded to the computing grid. When the intra-BAN protocol starts, all the sensors in one BAN send out probing packets while their CH collects the transmission power information from these sensors, selects the best transmission power levels, and broadcasts one command packet requesting the sensors to adjust their transmission powers so that energy consumption and generated interference can be minimized.

As our solution for prioritization of vital sign transmissions exploits channel diversity and LQI to avoid interference, we are interested in comparing it with:
- PROBE, which probes for the best channel based on the LQI and routes with the maximum route quality indicator (defined as the minimum LQI value among the links along a route);
- LEAST, which uses only one channel but does not rely on LQI and routes on the path with the least number of hops.

The curves for our solution in the following figures are denoted by X-MQMAC-CQBR-HS, where X represents the type of patient's vital sign data (i.e., whether it is ECG, CO2, or cardio-venous pressure [CVP] in that order of priority). As for e2e delay, our experiment results in Fig. 4a show that MQMAC-CQBR-HS outperforms both PROBE and LEAST. The hybrid scheduler (HS) in MQMAC-CQBR-HS guarantees delay performance; consequently, the packet error rate (PER) is reduced by exploiting channel diversity to allow simultaneous transmissions and avoid interference as well as by selecting the best route with best link quality. As far as e2e reliability (packet delivery ratio) is concerned, Fig. 4b shows that as the number of BANs increases, all these protocols experience decreasing reliability due to an increased interference and a higher number of hops to the sink. Due to similar reasons as above, MQMAC-CQBR-HS has the highest reliability while LEAST has the worst.

Context Awareness through Activity Recognition

As mentioned earlier, the success of ubiquitous healthcare systems lies in how intelligently these systems can take the user's context into account. Recognizing the user's current actions will help analyze the behavioral pattern of the user by correlating it with his/her environment and will also help determine the physiological state of the user by correlating it with processed vital signs. Sensor-based motion recognition, which uses data from wirelessly connected kinematic sensors attached to different body sites, integrates the area of sensor networks with novel data mining and machine-learning techniques to model a wide range of human motions. It enables a variety of applications such as rehabilitation, sports science/medicine, geriatric care, and health/fitness monitoring. Prior research on sensor networks concerned itself with data-centric self-configuration of network elements through low-level in-network processing of data. However, it largely overlooked the need for integration between self-configuration and data interpretation to enable intelligent sensor-based systems such as the human motion recognition systems.

In [5], we propose a novel window-based

Case Study	Actual Activity	No. Tests	No. Incorrect	Accuracy (%)
Case 1	Walking	12	0	100
	Standing	12	0	100
	Writing	12	0	100
	Smoking	12	0	100
	Jacks	12	0	100
	Jogging	12	0	100
Case 2	Walking	12	1	91.6
	Standing	12	0	100
	Writing	12	0	100
	Smoking	12	1	91.6
	Jacks	12	0	100
	Jogging	12	1	91.6
Case 3	Walking	24	2	91.6
	Standing	24	0	100
	Writing	24	0	100
	Smoking	24	2	91.6
	Jacks	24	0	100
	Jogging	24	1	95.8
Summary		288	8	97.2

Table 1. *Activity recognition algorithm performance.*

algorithm that can be tuned to recognize on the fly either various activities or movements in a specific activity using a supervised learning approach based on support vector machines (SVMs). To recognize various activities and their constituent movements, raw kinematic data from accelerometers (which provide triaxial acceleration) and gyroscopes (which provide angular velocities) needs to be collected for each activity and movement type considered. However, using all the raw data would be inefficient in addition to adding complexity. Hence, meaningful features need to be extracted from the raw data. We used features such as mean, standard deviation, maximum, peak-to-peak, root-mean-square, and correlation between pair of accelerometer and gyroscope axes. The features were then used to train a SVM to the type of activities or movements that need to be recognized.

When features extracted from a new set of data points are fed into the SVM, the algorithm solves a multi-class classification problem to recognize the activity or movement being performed. Our window-based algorithm can not only recognize on the fly an activity (or the movements involved in a particular activity) but also find the starting and finish instants of the current activity (or the movements in it). Our approach for recognizing activities (and/or the movements) consists of three phases — *training* (as SVMs are supervised learning machines), *tuning* (as the best set of features needs to be identified for improving the accuracy of recognition), and, finally, *real-time motion recognition*.

Sensor-based activity recognition plays a key role in the formation of context awareness, which involves the following three steps, namely, *perception*, i.e., identification and representation of the physical activities over time, *comprehension*, i.e., correlation of the identified activity with the vital sign data for detailed analysis, and *projection*, i.e., prediction of future vital sign trends to identify anomalies and comment on the health and wellness of a subject. One of the fundamental concerns in the theory of context awareness [2] is the problem of context identification, i.e., knowing which context is relevant to capture. This is addressed in our solution by training the SVM for only a specific set of activities we are interested in monitoring.

To show the performance of the window-based algorithm in recognizing various activities, we considered six different activities. For our study, we used Shimmer[5] biomedical sensor nodes for collecting tri-axial linear acceleration and angular velocities. The motes were placed on the right arm wrist and right foot of the subject. Each of these nodes gathered raw data at 20 Hz and transmitted it wirelessly to a destination node connected to a desktop that aggregated all the samples and extracted the features. We asked three subjects to perform all the six activities in the order they preferred. To evaluate the performance of activity recognition we considered three case studies:

- Train and recognize activities of each subject one by one;
- Train for activities of all subjects at one and recognize the activities of each subject;
- Train for one subject and recognize the activities of another subject.

As shown in Table 1, the overall accuracy of the system in recognizing the activities out of a total of 288 test was found to be 97.2 percent. Although our experiments were carried out in a controlled environment in which subjects were asked to perform repeatedly specific activities, the results of case studies 2 and 3 show that our method is also effective in less controlled environments, i.e., like in field trials (when given a larger training set that includes a wider range of movements and data from a large number of subjects).

Conclusion

We envision that future ubiquitous healthcare systems will be characterized by

- Pervasive vital sign monitoring using non-invasive sensors;
- Real-time processing of monitored data to derive meaningful physiological parameters;
- Context-aware data- and patient-centric decision-making.

[5] *Shimmer research: http://www.shimmer-research.com/*

We proposed an autonomic resource provisioning framework that can harness under-utilized computing resources in the vicinity to support real-time processing of vital signs (using inherently complex physiological models) and to acquire context awareness (using machine-learning-based algorithms). We also presented an innovative solution for reliable and continuous wireless transmission of vital signs as well as a window-based algorithm for acquiring context awareness using data from kinematic sensors.

Our ideas go beyond the specific application scenario presented in this article and can be applied to provide better healthcare and intervention in developing countries. For instance, a few wearable suits with non-invasive sensors can be time shared among inhabitants in densely populated villages or each person can be provided with one wearable suit in sparsely populated villages. The data can be locally stored on the sensor nodes and can be analyzed offline later using powerful processing devices such as a desktop computer in a delay-tolerant fashion. Alerts and reports of early diagnosis can be made available to health workers through a remote online electronic health record database. Analysis of stored data from multiple subjects under geographic and demographic context may reveal patterns of endemics/epidemics and enable timely medical intervention.

Alerts and reports of early diagnosis can be made available to health workers through a remote online electronic health record database. Analysis of stored data from multiple subjects under geographic and demographic context may reveal patterns of endemics/epidemics and enable timely medical intervention.

References

[1] S. Armstrong, "Wireless Connectivity for Health and Sports Monitoring: A Review," *British Journal of Sports Medicine*, vol. 41, no. 5, 2007, pp. 845–49.

[2] A. Talaei-Khoei *et al.*, "A Framework for Awareness Maintenance," *J. Network and Computer Applications*, vol. 35, no. 1, Jan. 2012, pp. 199–210.

[3] M. Satyanarayanan *et al.*, "The Case for VM-Based Cloudlets in Mobile Computing," *IEEE Pervasive Computing*, vol. 8, no. 4, Oct.–Dec. 2009.

[4] B. Chen and D. Pompili, "Transmission of Patient Vital Signs Using Wireless Body Area Networks," *Mobile Networks and Applications* (Springer), July 2010, pp. 1–20.

[5] J. P. Varkey, D. Pompili, and T. Walls, "Human Motion Recognition Using a Wireless Sensor-based Wearable System," *Personal and Ubiquitous Computing* (Springer), Sep. 2011, pp. 1–14.

[6] D. Estrin and I. Sim, "Open mHealth Architecture: An Engine for Health Care Innovation," *Science*, vol. 330, no. 6005, Nov. 2010, pp. 759–60.

[7] E. B. Lerner and R. M. Moscati, "The Golden Hour: Scientific Fact or Medical Urban Legend?," *Academic Emergency Medicine*, vol. 8, no. 7, July 2001, pp. 758–60.

[8] J. K.-J. Li, *Dynamics of the Vascular System*, World Scientific Publishing, Singapore, 2004.

[9] F. Dabiri *et al.*, "Ubiquitous Personal Assistive System for Neuropathy," *Proc. Int'l. Wksp. Systems and Networking Support for Health Care and Assisted Living Environments (HealthNet)*, Breckenridge, CO, June 2008.

[10] D. Malan *et al.*, "CodeBlue: An Ad Hoc Sensor Network Infrastructure for Emergency Medical Care," *Proc. Int'l. Wksp. Wearable and Implantable Body Sensor Networks (BSN)*, London, UK, Apr. 2004.

Biographies

Hariharasudhan Viswanathan (hari_viswanathan@cac.rutgers.edu) started his Ph.D. program at the Department of Electrical and Computer Engineering (ECE), Rutgers University, in 2009. Currently, he is pursuing research in the fields of mobile grid computing, data center management, and cognitive radio emergency networking under the guidance of Dr. Dario Pompili at the NSF Cloud and Autonomic Computing Center (CAC). Previously, he had received his B.S. in ECE from PSG College of Technology, India, and M.S. in ECE from Rutgers University in 2006 and 2009, respectively.

Baozhi Chen (baozhi_chen@cac.rutgers.edu) started his Ph.D. program at the Department of Electrical and Computer Engineering (ECE), Rutgers University, in 2007. Currently, he is working in the NSF Cloud and Autonomic Computing Center (CAC) under the guidance of Dr. Dario Pompili. His research interests lie in underwater acoustic communications and networking as well as wireless ad hoc and body-area networking with emphasis on cross-layer optimization. Previously, he had received his B.S. degree from Beijing University of Posts and Telecommunications, China, in 2000 and his M.S. degree from Columbia University, New York, in 2003.

Dario Pompili (pompili@cac.rutgers.edu) joined the Department of Electrical and Computer Engineering (ECE) at Rutgers University as an assistant professor in Fall 2007. He received his Ph.D. in ECE from the Georgia Institute of Technology in June 2007. He had previously received his Laurea (integrated B.S. and M.S.) and doctorate degrees in telecommunications and system engineering from the University of Rome "La Sapienza," Italy, in 2001 and 2004, respectively. In 2011, he won the NSF CAREER award for his work on underwater multimedia acoustic communication. In 2012, he won the ONR YIP award for his work on real-time in-situ data processing in heterogeneous mobile computing grids.