Cambridge University Press & Assessment
978-1-107-02102-0 — Body Area Networks
Safety, Security, and Sustainability
Sandeep K. S. Gupta , Tridib Mukherjee , Krishna Kumar Venkatasubramanian
Excerpt
More Information

1 Introduction

A healthy populace is essential for societal prosperity and well-being. This makes healthcare a basic societal necessity. Most societies endeavor to provide it in some form, usually in an organized manner through a system of medical facilities that are either private or public, or both. The current mode of delivering care relies on the patient initiating the care-delivery process. Figure 1.1 illustrates this model. It is a four-step process and requires the patient to observe the presence of specific symptoms, and visit a caregiver, who then diagnoses the problem and provides a treatment. We call this the *traditional model of care delivery*. The principal characteristic of the traditional model is that it is reactive in nature. No action is taken to improve the patient's condition unless the patient initiates the process. A problem with this approach is that it is inherently defensive in nature in fighting illness. This is particularly problematic if the symptoms for the patient's ailments seem benign or manifest late in the progression of the disease.

Additionally, the traditional model has other problems as well, especially those associated with scale. The whole model of care was designed for a society where the number of people requiring care was a very small percentage of the population. However, with the dramatic demographic shift taking place in the world, especially that associated with aging, the traditional model of care delivery will be a huge bottleneck. Furthermore, the societal demographic shift is also producing one at the level of caregivers, leading to a dramatic decrease in the number of caregivers available in certain specialties [1]. This will most likely lead to dire shortages of healthcare personnel, and, if this situation is left unattended, could result in a drop in the quality of medical care and a substantial increase in healthcare costs [2, 3]. In summary, with the traditional model of care delivery in place, the healthcare system in most countries will (the US healthcare spending is projected by the Centers of Medicare and Medicaid to be 4.5 trillion by 2019 [4]) increasingly come under pressure as the average age of their population increases and the number of elderly people swells.

One of the ways of reversing this trend is by the introduction of automated pervasive health-monitoring technologies that can monitor a person's health and alert appropriate healthcare personnel in case of emergencies, thereby providing optimal care with minimal supervision. In response to this need, embedded computing devices such as smart phones, iPads®, and other personal electronic devices are being increasingly used to support health-related applications such as tracking calorie intake [5], weight control [6], pulse oximetry [7], and educational training for nurses [8]. Trends indicate a fast progression towards pervasive health-monitoring systems (PHMSes), where sensors

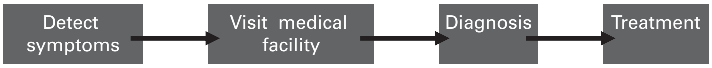

Figure 1.1 The traditional healthcare-delivery model.

are deployed on the human body to monitor environmental and physiological signals, a smart phone acts as a computation and communication hub, or base station, for these sensors, which gathers data and processes them to detect contexts, such as location, activities, or health emergencies, and a personal laptop uploads information to the cloud to form a history of health. With the smart phone being recently considered as a medical device [9], PHMSes are predicted to become critical infrastructures to support the health and well-being of the populace. *Pervasive healthcare* can play a major role in providing continuous care to patients, thereby altering the traditional care-delivery model into one that provides healthcare facilities to individuals anywhere and at any time. It uses large-scale deployment of sensing and communication (wired and wireless) technologies to monitor patients continuously. This allows it to deliver accurate health information to caregivers, thereby stimulating timely diagnosis and treatment for health problems.

1.1 Pervasive healthcare

Significant advances in communication and sensing technologies have led to the development of intelligent hand-held and wearable devices (such as smart phones, smart clothes, and smart apparel). Such devices provide a platform to implement a wide range of solutions for health-monitoring purposes. Examples include vital-signs-monitoring apps available on various smart phones. Further, their wearable or hand-held nature allows them to be present around the user at all times, making health monitoring a pervasive activity. We call such technologies that use pervasive computing capability for health management *pervasive healthcare systems*.

The health-monitoring capability of pervasive healthcare systems makes them ideal for many diverse applications [10], including the following. (1) *Telemedicine*. This provides the ability to monitor, diagnose, and treat patients anytime and anywhere. The automation brought about by the technology further reduces the chances of errors and enables timely treatment of patients by providing accurate, real-time, and complete health information to the caregivers. (2) *Lifestyle management*. Pervasive healthcare systems have the ability to provide personalized care by keeping track of a user's health in minute detail and providing the information as needed. For example, it can be used by people to improve their health by developing specialized meal and exercise plans. (3) *Data repositories*. Pervasive healthcare systems are designed to collect data from patients over long periods of time. These data are stored in an organized manner so that they can be studied by the patients' caregivers in order to provide better care. Such large data sets can be useful for studying issues such as the response to medicine, the demographics of people with specific ailments, possible improvements in the care, improvement in medicine, alternative treatments, and diagnosis.

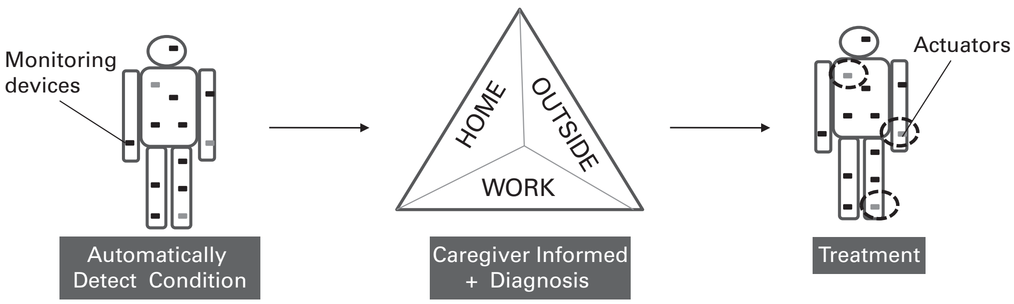

Figure 1.2 The envisioned healthcare-delivery model.

Pervasive healthcare systems by their very nature alter the traditional treatment model by making it more proactive. They provide a proactive patient-centric form of care (as opposed to the reactive caregiver-centric form of today's care), which permits health problems to be detected and responded to as soon as they occur. Now, the presence of an ailment will be detected way in advance before some or even any of the symptoms manifest themselves. This detection of the ailment will cause the information to be made available to the caregiver directly. No personal visit is required immediately. The caregiver analyzes the data, performs diagnosis, and specifies treatment for the user, which can be provided to the user through the use of actuation devices or may require the patient to visit the caregiver. An additional advantage of the pervasive monitoring system from the perspective of the caregivers is that a single caregiver can now manage a large number of patients. Further, this can be done without any time or space restrictions for either the caregiver or the patient. Figure 1.2 illustrates this newly envisioned healthcare-delivery model.

Conceptually, the model consists of three main planes. The model has been adapted from the one presented in [11]. Figure 1.3 illustrates the three planes. The three planes are the device plane, the data-access plane, and the knowledge plane. The *device plane* is used to collect medical data from patients and perform preliminary processing on it. It consists of a large number of medical devices that possess the capability to collect both physiological and ambient data from around the patients. The devices may be deployed on patients in both in-vivo and in-vitro manner. The data collected from the devices are handled by the *data-access plane*, which serves three purposes: (1) enabling the devices to interact seamlessly with one another; (2) providing facilities for organizing pre-processed health data from devices into a structured electronic format (also known as electronic patient records (EPRs), electronic health records (EHRs), or electronic medical records (EMRs)), and for their long-term storage; and (3) enabling device actuation as a treatment measure. The data-access plane also provides the infrastructure for managing the health data collected by the devices. Finally, the *knowledge plane* is used for reasoning on the data in the EPRs. Some of the features it can provide include detection of the occurrence of a medical emergency, the failure of a specific treatment procedure, and inconsistencies between the proposed diagnosis and the symptoms. This capability gives the caregivers feedback pertaining to their diagnoses and treatment,

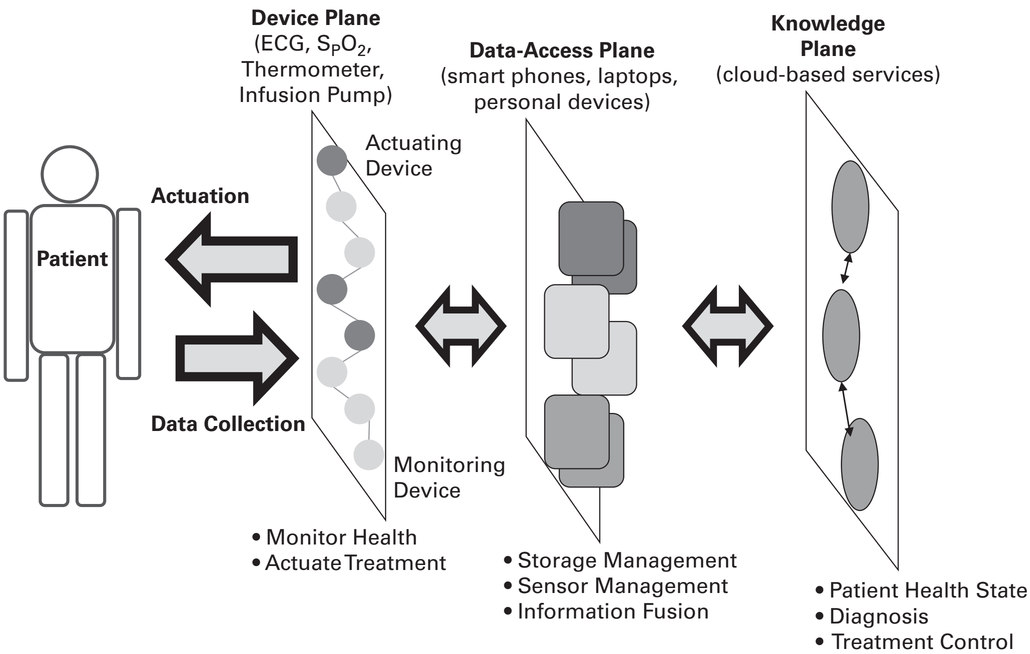

Figure 1.3 The pervasive healthcare model: conceptual view.

allowing them to make appropriate adjustments through the data-access plane. Recent initiatives by the US government have resulted in the Health Information Exchange (HIE) cooperative agreement program (`http://healthit.hhs.gov/portal/server.pt?open=512&objID=1488&mode=2`), which encourages exchange of health information between different healthcare facilities throughout the USA.

1.2 Monitoring technologies

One of the most pressing needs in modern medicine is to develop medical devices capable of providing continuous care (i.e., monitoring, decision support, and delivery of therapy). Such devices are expected to decrease healthcare costs by enabling alternatives such as home-based or ambulatory care. This allows caregivers to permanently have a detailed picture of the patient's health, enabling them to better tune the treatment provided. Such a system also makes possible real-time notification in the event of emergencies and providing first-responders with accurate and complete information about the patient's health.

These new-generation medical devices are usually cyber-physical in nature; that is, they are very tightly coupled with their environment (the human body). Further, since they are being designed for long-term monitoring purposes, they usually have a very small form factor with minimal power footprint, and use wireless communication technologies where possible. They have been designed to monitor a plethora of ailments,

Table 1.1 Prominent pervasive care systems

Pervasive care system	Sensor type	Purpose
SenseWear[a]	Wearable monolithic device	Physical-activity and sleep-behavior monitoring
Microcompaq[b]	Wearable monolithic device	Continuous monitoring of vital signs
UbiMon[c]	Wearable and implantable sensors	Capturing transient but life-threatening medical events
HealthGear[d]	Wearable sensors	Detecting sleep apnea
MyHeart[e]	Biomedical smart clothes	Monitoring/detecting cardiovascular diseases
CodeBlue[f]	Wearable sensors	Emergency care, disaster response, and stroke rehabilitation
Lifeshirt[g]	Wearable sensors	Continuous monitoring of vital signs
Ayushman[h]	Wearable sensors	Continuous monitoring of vital signs

[a] http://sensewear.bodymedia.com/
[b] http://www.welchallyn.com/
[c] http://www.doc.ic.ac.uk/vip/ubimon/home/index.html
[d] http://research.microsoft.com/~nuria/healthgear/healthgear.htm
[e] http://www.hitech-projects.com/euprojects/myheart/
[f] http://www.eecs.harvard.edu/~mdw/proj/codeblue/
[g] http://www.lifeshirt.com
[h] http://impact.asu.edu/Ayushman.html

such as cardiovascular diseases and neurological problems, rehabilitation, and general health, and for collecting meta-physiological-state information (sleep, awake, fatigue), circadian activity monitoring, extreme-environment medical monitoring (e.g., space), and sleep monitoring. The biomedical platforms developed in this regard can be broadly classified into three broad categories: (1) monolithic platform-based, (2) textile (fabric)-based, and (3) body-sensor-network-based. Some of the prominent health-monitoring BANs are listed in Table 1.1.

Monolithic platform-based solutions

Monolithic solutions consist of devices that are basically a micro-controller board used as a medical monitoring platform. Various (often custom-designed) physiological monitoring sensors (biosensors) are used to collect specific physiological parameters and transfer them to the board for signal processing and storage. The monitoring platform usually forwards the patient data to the back-end (hospital network) for storage. This is usually done via a wireless channel in order to ensure the patient's mobility. Such monolithic monitoring devices allow the monitoring of a diverse set of physiological signals such as the electrocardiogram (ECG), photoplethysmogram (PPG), electromyogram (EMG), galvanic skin response (GSR), skin temperature, and blood pressure (BP). Some of the prominent projects in this category include the LiveNet at MIT [12], AMON at the ETHZ [13], and LifeGuard at NASA Ames and Stanford [14]. Such systems, given their monolithic nature, are easy to manage; however, they possess minimal intelligence, and packing all the sensors with one board makes the entire setup bulky, with relatively large form factors affecting its usability.

Textile-based solutions

Textile-based medical monitoring platforms try to alleviate some of these problems by distributing the sensors over a fabric worn by a person, making them truly wearable health-monitoring systems. Textile-based platforms integrate biosensors into the fabric of the garment by using conductive fibers, which are knitted like regular yarn. The entire system is supplied by a centralized portable power source, which is a part of a micro-controller board that is used to collect, process, and possibly forward (to back-end entities, wirelessly) data from the sensors. No wireless module is required at the sensors since the fabric itself acts as a data-bus. As in the previous case, textile-based monitoring platforms enable the monitoring of a slew of physiological parameters such as the ECG, EMG, activity (accelerometer), PPG, skin temperature, GSR, BP, and body temperature. Some of them even have the ability to raise alarms if the physiological parameters go beyond preset thresholds [15]. The platforms typically differ in terms of the fibers used for knitting the garment (conductive and/or electrostrictive), the type of garment (vest or covering the entire upper body), and usability (washable, stretchable). Some of the prominent systems which use the textile-based approach to monitoring include MyHeart [16], Wealthy [15, 17, 18], MagIC [19], and SmartVest [20]. The systems also vary in terms of the available signal-processing capabilities, with respect to removing measurement artifacts, removing baseline noise from ECG signals, and individualized calibration of sensors. Textile-based systems have the advantages of being easier to use, and providing good signal quality and comfort. However, the dramatic shift in paradigm has made their adoption very slow. In [21], the authors posit that the principal reasons for this are the limited involvement from the stake-holders, the yet-to-be-proven clinical effectiveness of the technology, and the poor cost-to-benefit ratio for payers (insurance agencies) thus far, in terms of safety, security, and convenience.

Body-area-network-based solutions

The third category of next-generation medical devices are those which form a network on the person being monitored using a large number of low-capability computing platforms called sensors. These body sensor networks (BSNs) or body area networks (BANs) use tiny wireless-enabled processing units (often referred to as motes in the literature), which are usually interfaced with one or more custom-designed or off-the-shelf sensing modules, also called *motes*. The data from the sensing modules are collected by the mote, and forwarded wirelessly to a base station for long-term storage and processing. The motes equipped with a micro-controller can do limited amounts of processing, such as reducing motion artifacts and carrying out signal cleaning. The wireless channel used also varies with the platform used, with ZigBee, IEEE 802.15.4, and Bluetooth being the prominent technologies used. Some of the prominent applications based on BANs include CodeBlue at Harvard [22], the Alarm-Net platform at UVa [23], Human++ [24], BASUMA [25], and HeartToGo [26]. The BAN-based systems have the advantage of being easy to deploy and manage. However, they do suffer from poor signal quality due to motion artifacts and from poor battery life, and the wireless channel used is susceptible to eavesdropping and jamming.

1.3 Overview of the book

In this book, we focus on the third type of continuous monitoring technologies available, i.e., body area networks. The principal idea is to ensure that BANs do not cause potential hazard to the human body from interruptions in, and side-effects of, their operation. For example, in a cardiac-care scenario with a BAN made of heart-function monitors (ECG, PPG sensors) and actuators (pacemakers), the potential hazards of operation include battery exhaustion (this interrupts the BAN's operation), initiation of unauthorized and untimely actuation (shocks) by replaying old commands (i.e., triggering unauthorized actions), and unwarranted electromagnetic interference from nearby electronic devices such as mp3 players inducing effects on long-term BAN operations. Consequently, the focus of this book is on how BANs can be designed in order to make them more safe, secure, and energy-wise sustainable for long-term operation. The emphasis is placed on software-engineering aspects of BANs.

The software-design methodology that is considered in this book is that of *model-based engineering* as shown in Figure 1.4. In this methodology, the given BAN system and its interaction with the human body and the environment (discussed in Chapter 2) are first represented using mathematical abstractions or models (Chapter 3). These models can be behavioral models that represent the operation of the BAN using transfer functions (Chapters 4 and 6), formal models that represent the system using finite-state machines (Chapters 4 and 6), hybrid systems, which model the cooperation of the embedded systems with the human physiology (Chapters 4 and 6), or feature-based

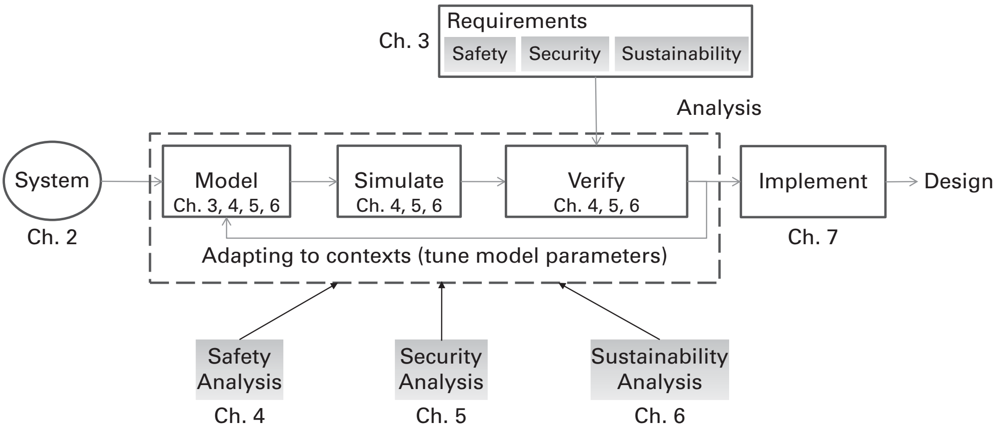

Figure 1.4 Model-based environmentally coupled solutions for BAN design.

representations obtained using signal-processing algorithms (Chapter 5). The models are then simulated, whereby the model parameters are mathematically manipulated in order to obtain the system behavior in terms of system properties. The system properties are then checked against the safety, security, and sustainability requirements (Chapter 3) to check for their violation. If there are violations, then the models are adapted to the environmental context, which includes the physiological state of the BAN user. The adaptation procedure typically involves tuning the model parameters or changing thresholds in the signal-processing algorithms. This procedure, also known as model-based analysis, is discussed in detail for each of the BAN requirements of safety (Chapter 4), security (Chapter 5), and sustainability (Chapter 6). The output of the model-based analysis phase is a BAN model that is verified against requirements. The next stage is the implementation of the model, which is handled by the design phase. The design phase is discussed in Chapter 7, where several automatic code-generation tools are described.

1.4 Questions

1. What are the advantages and disadvantages of the traditional healthcare model as described in Section 1?
2. Think of three potential applications for pervasive healthcare systems apart from the applications described in this chapter. Give as much detail as possible, including design diagrams.
3. One of the advantages of using pervasive healthcare systems is the automated population of electronic patient records (EPRs). These have slowly come to replace paper-based records in care facilities and afford many advantages. What are the advantages of EPR systems over paper-based records, and what, if any, are the potential disadvantages?
4. Imagine a hospice or elderly-care facility with about 10 patients being cared for. There is one medical team responsible for the health of all the patients, consisting of a primary-care provider, a geriatrics specialist, and a couple of nurses. Design a health-monitoring system for such a facility that can monitor patients continuously and provide the information to the care-providers using the technologies (one or more) described in this chapter. Discuss the potential performance, scalability, and usability of the system designed.

2 Body area networks

A body area network (BAN) is a network consisting of a heterogeneous set of *nodes* that can sense, actuate, compute, and communicate with each other through a multi-hop wireless channel. A BAN collects, processes, and stores physiological (such as electrocardiogram (ECG) and blood pressure), activity (such as walking, running, and sleeping), and environmental (such as ambient temperature, humidity, and presence of allergens) parameters from the host's body and its immediate surroundings; and can even actuate treatment (such as drug delivery), on the basis of the data collected. BANs can be very useful in assisting medical professionals to make informed decisions about the course of the patient's treatment by providing them with continuous information about the patient's condition.

BANs belong to a much more generic class of device networks called wireless sensor networks (WSNs) [27]. BANs evolved from WSNs through a series of intermediate steps whereby first the WSN concept was applied to personal devices such as laptops, phones, cameras, and printers. Such networks are called personal area networks (PANs) [28], or wireless PANs (WPANs) [28]. From WPANs evolved BANs in which medical devices, such as pulse oximeters, and personal computing or auxiliary devices such as smart phones and retina prostheses [29, 30] were networked through the wireless channel and worn on the body. Devices were also implanted, such as pacemakers, which communicate through the body to an outside controller. In a hospital setting, BANs are networked with other in-hospital medical devices such as Holter monitors and medical data recorders to form medical device networks (MDNs) [31] for post-operative or intensive-care-unit (ICU) patient monitoring.

Sensors form the essential basis of a BAN and come in different forms. Essentially sensors in BANs may have one or more of the following properties.

- They may be *heterogeneous* in terms of capabilities, and are designed to be *unobtrusive* to the host. Consequently, individual sensors in a BAN may have very *limited* form factor, power source, memory, computation, and communication capabilities, compared with generic sensor nodes, thus requiring BANs to employ a large number of nodes in order to collect patient health data in a reliable and fault-tolerant manner.
- They may be *implantable or wearable* in nature, implanted sensors are subject to more stringent requirements with respect to safety, security, and sustainability than do wearable ones.
- They may be designed to *measure multiple stimuli* from their environment (the human body).

- They may be designed to *survive extreme conditions* such as variation in temperature and the presence of water or saline [32].
- They may be *powered using scavenged energy* sources such as body movements, body heat production, flexible solar cells and biofuels such as blood glucose [32].

For example, implantable pacemakers need to operate in the presence of saline body fluids and may be operated using energy scavenged from body heat [33].

Each BAN has a controlling entity called the *base station*, which collects and processes data for the BAN. All the sensors in the BAN communicate the data they collect to the base station at regular intervals. BANs can vary in size from a network of a few nodes to a large network with a few hundred. Systems with upward of 150 sensors have been designed [34]. It has been suggested in [35] that a large number of low-quality sensors can perform the task of monitoring as effectively as a few high-quality sensors, while at the same time being less intrusive. Consequently, we assume that BANs with hundreds of sensors are feasible. Needless to say, BANs have many diverse applications, including sports health management, home-based healthcare for the elderly, and post-operative care.

2.1 Architecture

A BAN consists of a network of computing units that are used to perform context-aware ubiquitous operations to aid monitoring and autonomous actuation applications in healthcare, and for performance monitoring and giving feedback in military, leisure, and sports applications.

The nodes in a BAN are of varied capabilities; however, they can be broadly classified into two categories: (1) *sensor nodes*, which are implanted or wearable medical devices or simply low-capability wireless computing platforms interfaced with sensors and actuators (e.g., a PPG sensor interfaced with TelosB motes); and (2) *base-station* nodes, which have higher computational and communication capabilities (e.g., smart phones) than the worker nodes, and are used for disseminating information to and collecting information from the worker nodes. The base station controls the entire BAN and can reach every node in it in a single hop. It has significantly higher computational and communication capabilities than do the sensors, therefore all sensors send their data to it for processing. Note that, unlike in traditional wireless sensor network applications where there is limited control over the sensor networks after deployment, the BAN is a fairly controlled system, being under the supervision of the patient or medical personnel at all times. IEEE Task Group 6 (TG6) (`http://ieee802.org/15/pub/TG6.html`) has defined the standard architectures of a BAN. According to the standards the prominent characteristics of a BAN architecture can be summarized as follows.

2.1.1 Hardware

Sensors in a BAN are heterogeneous with respect to hardware configuration. With recent advances in embedded-systems research a plethora of small-scale computing