Model-Based Architecture Analysis for Wireless Healthcare

Amitabha Ghosh
Dept. of Electrical Engineering
Princeton University, NJ
amitabhg@princeton.edu

Ying-Kei Hui
Dept. of Internal Medicine
Pennsylvania Hospital, PA
ying-kei.hui@uphs.upenn.edu

Mung Chiang
Dept. of Electrical Engineering
Princeton University, NJ
chiangm@princeton.edu

ABSTRACT

The primary challenges in deploying a wireless healthcare solution stem from real-time, distributed resource constraints, as well as stringent clinical requirements of reliability, safety, device interoperability, QoS guarantee, and privacy/security. Although optimized solutions exist for each individual element of the system, the complex, distributed, and concurrent interactions among multiple subsystems make system integration a costly bottleneck. We show that end-to-end modeling and analysis using the formalisms of architecture description languages like AADL can alleviate the hurdles of system integration and provide an effective way of addressing these challenges, thereby, making deployments possible.

Categories and Subject Descriptors

D.2.8 [**Software Engineering**]: Metrics—*performance measures*; D.2.1 [**Software Engineering**]: Requirements / Specifications—*software architectures, interoperability*

General Terms

Design, Reliability, Performance

1. INTRODUCTION

Driven by recent technological advances in low-power medical sensors and wireless communications infrastructure, as well as the need for real-time continuous monitoring of patients' physiological data, networked wireless devices are envisioned to show tremendous promises for an improved future healthcare. Cost effective medical sensors can now be equipped with wireless connectivity, such as WiFi, Bluetooth, ZigBee, and near-field radios, using which they can communicate with smart phones, wireless sensors, personal digital assistants (PDAs), and bed-side computers. Wireless solutions, in addition to being cost effective for indoor hospital environments, have an equally great potential for telemedicine and biotelemetry, especially in first responder and disaster recovery scenarios [6].

MobileHealth'*11*, May 16, 2011, Paris, France.

However, despite this promising vision, we have seen only a few deployments of wireless solutions [5] to date in a clinical environment. The key challenges, perhaps, lie in the stringent medical requirements of reliability, fault tolerance, safety, device-interoperability, QoS guarantee, and, most importantly, privacy/security. Such requirements are further exacerbated by real-time resource constraints, as well as complex, distributed, and concurrent interactions among multiple, heterogeneous susbsystems used in a clinical environment. This calls for the need of *end-to-end system modeling* and *architectural analysis* prior to implementation [3,4].

Traditionally, software engineering practice advocates building and testing each component separately, while leaving their individual assumptions being stated informally and implicitly without regard to their mutual compatibility. This leads to serious faults and instability during system integration. To avoid such issues, in recent years, architecture description languages like AADL (Architecture Analysis and Design Language) [2] have seen an increasing use for end-to-end system modeling and analysis, which have proved to be significantly helpful in early discovering of anomalies and identifying root causes in mission-critical applications, such as in the field of avionics, aerospace, and automotive systems.

In wireless networked systems deployed in a clinical environment, where heterogeneous devices and protocols are expected to seamlessly interface with each other, often under real-time resource constraints, the use of AADL can be extremely useful. In particular, the *component-connector* paradigm of AADL not only provides a precise, machine-processable way of describing software and hardware architectures, but also facilitates end-to-end analysis for various performance metrics using tools like OSATE [1]. The complexity of integrating and verifying a wireless healthcare system arises because interaction bugs among heterogeneous devices and protocols are unlikely to show up in the early development phase, and can trigger second or third order faults in the late integration phase causing system crashes. Thus, early system validation by – *automatic model checking* and *implementation compliance verification* – can avoid implicit mismatched assumptions across components and prevent system instability.

In this on-going work, we explore the challenges and provide future research directions by examining two specific medical scenarios where wireless solutions can substantially improve the quality of healthcare, and advocate the use of AADL-inspired end-to-end system modeling and architectural analysis.

2. AADL OVERVIEW

AADL provides a rich set of modeling formalisms to describe architectural elements as a collection of interacting components, with their externally accessible features and configurations, and internal implementations. This involves formalization of the detailed requirments as an accumulated set of properties, followed by reasoning about design choices. This allows constructing a model of the entire system as well as ensuring interoperability among sub-system components [2].

Once a model is constructed, its behavior can be simulated to mimic the interactions between appropriate components of the system, as well as analyzed with realistic workloads using AADL toolsets and plug-ins. For wireless healthcare, the choice of the interacting components depends on the underlying protocols, network infrastructure (e.g., celluar, WiFi, low-rate wireless personal area networks), and the specific application. Therefore, our first task is to identify the key elements of the system and describe them using appropriate AADL components by defining their configurable properties. Then we need to instantiate each component with specific property-values and annotate their dependencies with other components. This flexibility allows users to predefine components so that different settings can be investigated during analysis.

3. WIRELESS HEALTHCARE SCENARIOS

In this section, we describe two clinical scenarios where wireless solutions can be of great use, and highlight the underlying challenges in implementation. We then construct AADL models, using which an end-to-end architectural analysis can address these challenges, thereby, making implementation and deployment possible.

3.1 In-Hospital Monitoring

Consider a hospital environment where patients with chronic or critical illness conditions are being continuously monitored in real-time for their vital signs. In particular, imagine a *telemtery monitoring* scenario where ECG (Electrocardiogram) signals (heart rate and rhythm), blood pressure, glucose level, respiratory rate, oxygen saturation level (i.e., pulse oximetry), and temperature are measured continuously using medical sensors from different groups of patients located on different floors in the hospital, or during night shifts when fewer medical staff are in-house. These measurements are then displayed on patient bed-side monitors, as well as transmitted reliably and securely over a multihop wireless network to a centralized station, where nurses and doctors can watch for potential emergencies and respond in a timely manner. As simple as this setup sounds like, the challenges in actually integrating the components and implementing the overall system are many fold.

3.1.1 Key Challenges

– *Deployment and Resource Management*: It is very likely that in a hospital environemnt, many other wireless devices operating on the same or adjacent frequencies are present, which can cause serious interference with the ones deployed for medical monitoring. For instance, smart phones operating on WiFi, microwave ovens, and cordless phones are known for causing strong interference. In addition, wireless signals get severely attenuated due to obstacles, such as walls and furniture, in an indoor environment. To overcome these issues, a planned deployment of the wireless infrastructure as well as designing algorithms for distributed allocation of shared resources (e.g., bandwidth, frequency, scheduling) to minimize interference is needed.

– *Wireless Infrastructure*: What kind of wireless infrastructure is suitable for such in-hospital, real-time transmission of sensitive medical information? For instance, is a network of cheap, low-power, battery operated wireless sensors, which can be placed at strategic locations inside the hospital building and can transmit over the unlicensed ISM (Industrial, Scientific, and Medical) band at low data-rates good enough? or, is a combination of Bluetooth devices, laptops, WiFi access points, as well as low-power wireless sensors needed? If so, what are their interoperability issues?

– *Security*: If the wireless infrastructure comprises low-power, battery operated sensors, then what kind of encryption algorithms can be used to prevent security attacks, and how much resources (e.g., CPU time, memory, energy) will their operations consume? In addition, given that wireless sensors can randomly fail, how do we ensure that mallicious attackers are not able to exploit such vulnerabilities? On the other hand, if the infrastructure has heterogeneous wireless devices, then agreeing on a common security protocol is the key challenge. To this end, designing light-weight security protocols that can work across device boundaries under real-time resource constraints is needed.

– *Quality of Service*: Depending on the seriousness of medical conditions, how should priorities be assigned to the corresponding traffic to ensure timely delivery of important packets, especially in the face of severe resource constraints? For instance, ECG signals are vital to a patient's heart condition and can cause death if not addressed immediately, as opposed to, say, blood pressure. Guaranteeing end-to-end latency requirements even in the presence of unreliable, interference-limited channel conditions and node failures is a major challenge.

– *Standardizing Interfaces*: There is a lack of standardized interfaces and protocols using which medical sensors attached to a patient's body can communicate with wireless devices that are capable of transmitting signals over multiple hops.

3.1.2 AADL Architecture Description

At a high level, we identify four key elements in this wireless healthcare application: (1) medical sensors, (2) wireless nodes, (3) a network, and (4) a wireless channel. We model the medical sensors as *device* components, the wireless nodes and the network as composite *system* components, and the wireless channel as a *bus* component. A wireless node, in turn, consists of four elements: a radio, a CPU (with processor and memory), a battery, and a GPS. Each of these can be modeled using appropriate AADL components. The network component is composed of three elements: a physical layer, a network layer, and a data-link layer, which includes a security and a MAC layer. We model each of these as a *system* component. The network layer component is responsible for multihop routing, whereas the security, MAC, and physical layer components together form the basis of secure and reliable hop-by-hop delivery.

We now choose configurable properties for these components, as shown in Figure 1 using AADL graphical notations. Although the exact list of properties for a given component depends on the specific instance chosen, in most cases, there

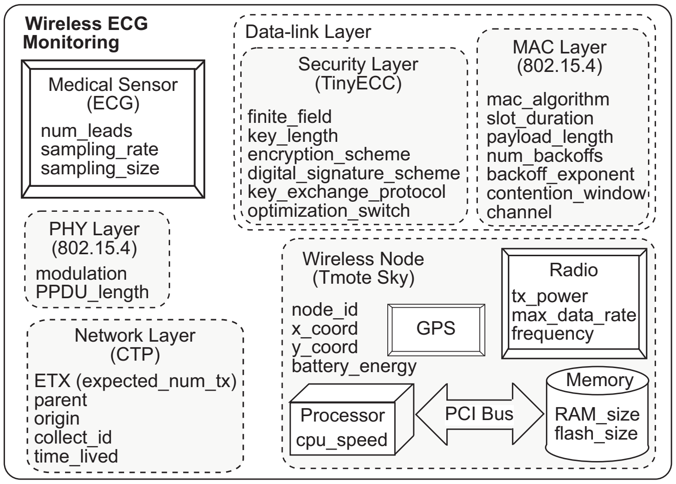

Figure 1: Configurable properties and specific component instances of the Wireless ECG Monitoring app.

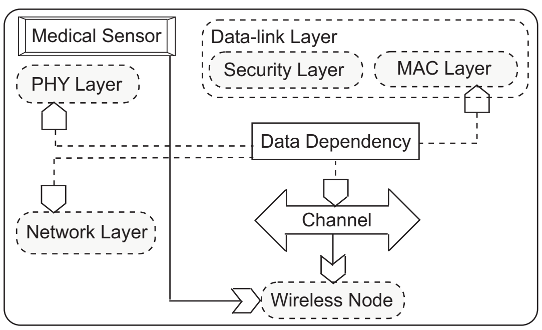

Figure 2: An instance of Wireless ECG Monitoring.

is a set of common properties that depends only on the functionality of the component. For instance, a public key infrastructure (PKI)-based security protocol will have multiple keys, an encryption algorithm, and a signature scheme, whereas a contention-aware MAC protocol will have a contention window, number of backoffs, slot duration, etc.

In our model, we chose an ECG machine that measures electrical activity of the heart muscle using leads as an instance of a medical sensor. For wireless nodes, we choose low-power Tmote Sky motes running TinyOS platform with TI MSP430 microcontroller. The address-free Collection Tree Protocol (CTP), which generates routes using a gradient-based approach, is used in the network layer. The security layer is instantiated with TinyECC (a type of PKC based on elliptic curves and supported by TinyOS), and the MAC and physical layers are modeled using IEEE 802.15.4 (ZigBee). Once the specific instances are chosen, the property-values are filled up, which can be used during analysis.

As a final modeling step, we use the AADL *Data Accesses* to represent the dependencies across components so that dependent property values can be coordinated. For instance, the data rate of the physical layer not only depends on the modulation scheme and the PPDU length, but also on slot duration, which is a property of the MAC layer. In Figure 2, we show an instance of the complete model using graphical AADL syntax.

3.2 Biotelemetry in Disaster Recovery

Biotelemetry is defined as transmitting physiological data to a remote location for further interpretation and decision-making. Consider an earth-quake affected environment where first responders are attending to hundreds of injured people who have been severely affected. Many of these people are elderly or with physical disabilities, who are stuck inside buildings and are so badly injured that they cannot be transferred until they receive life-saving medical treatment or their conditions are stabilized. However, since first responders are typically not doctors or nursing staff, and the medical equipments available to them are limited, such an emergency situation calls for external medical help from healthcare professionals in nearby hospitals, so that tailor-made instructions for patients can be delivered to the scene.

Smart phones, which are carried by almost every individual these days and come with various built-in sensors that can act as a stethoscope, including heart beat, bowel sound, and respiratory information, can be of great use – they can capture images and videos of wounds on a victim's body, as well as collect physiological data and transmit (or continuously stream) them securely along with patient's medical history over a cellular network. In disaster recovery scenarios where time is of essence, the use of such wireless devices can save lives.

However, as in the case of in-hospital monitoring, we face some similar challenges related to content-aware resource management, context-aware privacy, and security. In Table 1, the stars represent varying degrees of relative importance among traffic volume, sensitivity to errors, to delay, and to security concern, for different types of medical content. Such understanding and classification of priorities across the metrics is important for allocating limited resources in the paradigm of *content-aware* networking, QoS, and privacy/security.

3.2.1 Key Challenges

Content-Aware Networking: Current 3G cellular technologies, such as CDMA2000 and its EV-DO feature, can theoretically support a forward link rate of up to 2.4 Mbps and a reverse link rate of up to 1.8 Mbps. However, these data rates are unlikely to support simultaneous uplink transmissions of multiple real-time video streams from closely located and interference-limited mobiles. In addition, current cellular infrastructures do not have provisions for treating clinical video traffic differently, nor do they support content-aware resource sharing. For instance, the network can adapt itself by choosing optimal transmission power levels for individual video frames (instead of blasting at maximum power), or by selectively dropping unimportant frames (e.g., B frames) so as to fit multiple video streams within the available bit rate. This underlines the need for designing content-aware network protocols which can support multiple levels of quality of service.

Context-Aware Privacy: First responders in a disaster recovery scenario need access to patients' medical records. However, privacy policies are usually specified in human languages and lack the necessary clarity. Thus, there is a need to formally define context-aware privacy policies depending on patient-profiles, location, timestamp, and physiological parameters that can be reverse-engineered into patient IDs, which can then be processed by a machine and enforced while transmitting sensitive patient data.

Table 1: Clinical content and their properties

	Privacy Sensitivity	Delay Sensitivity	Error Sensitivity	Data Volume
Vital signs	***	***	***	*
ECG	***	***	***	***
Heart sound	*	**	*	**
Respiratory sound	**	**	***	**
Abdominal exam	**	**	***	**
Chest X-ray image	**	*	**	**
Ultrasound video	***	*	***	***

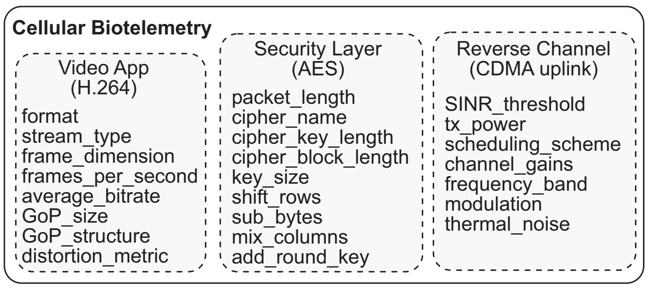

Figure 3: Properties of Cellular Biotelemetry.

Light-Weight Security: The encryption algorithm used in CDMA2000 is the Advanced Encryption Standard (AES) with 128-bit key that provides a very high level of security. For a disaster recovery scenario, the important question to ask is how long would AES take to encrypt the frames of a real-time video traffic, and whether it could still support the stringent latency bound. If not, designing new light-weight security protocols, especially suited for clinical wireless applications, is necessary.

3.2.2 AADL Architecture Description

We identify five key elements in this cellular technology based medical application: (1) a cell phone, comprising a built-in camera and a video app, (2) a base station, (3) a network, (4) a forward channel, and (5) a reverse channel. As in the case of in-hospital monitoring, each of these elements can in turn be composed of multiple sub-elements, which can be modeled using appropriate AADL abstractions, along with their configurable properties and values according to CDMA2000 specification. In Figure 3, we only show the properties of the video app, the security layer, and the reverse channel, which are modeled as *system* components.

4. CONCLUSIONS AND FUTURE WORK

In this on-going work, we have advocated the use of AADL in building end-to-end system model and doing architectural analysis that can help in overcoming the challenges of implementing a wireless healthcare system. We believe such formalisms are useful in system integration of real-time, complex, distributed application such as wireless healthcare. Our future work lies in using these models and AADL toolsets for end-to-end latency, security, and video distortion analysis.

5. REFERENCES

[1] P. H. Feiler, A. Greenhouse, and L. Wrage. Open Source AADL Tool Environment (OSATE) Plug-in Development Series. *http://www.aadl.info, CMU*, December 2005.

[2] P. H. Feiler, B. Lewis, and S. Vestal. The SAE Architecture Analysis and Design Language (AADL) Standard: A Basis for Model Based Architecture Driven Embedded Systems Engineering. In *RTAS*, May 2003, pp. 1–10.

[3] A. Ghosh, L. Pereira, T. Yan, and H. Cao. Modeling Wireless Sensor Network Architectures using AADL. In *ERTS*, Toulouse, France, January 2008.

[4] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, and P. Kristofer. System Architecture Directions for Networked Sensors. In *SIGPLAN Not.*, 35(11):93–104, November 2000.

[5] J. Ko, *et al.* MEDiSN: Medical emergency detection in sensor networks. In *ACM Trans. on Embed. Comput. Syst.*, 10(1):1–29, August 2010.

[6] D. Malan, F-J. Thaddeus, M. Welsh, S. Moulton. CodeBlue: An Ad Hoc Sensor Network Infrastructure for Emergency Medical Care. In *MobiSys WAMES Workshop*, April 2004.