

Research Statement

Sye Loong Keoh
sye.loong.keoh@gmail.com

December 23, 2009

1 Research Overview

My current research interests include security and privacy, policy-based management, trust management, middleware for wireless ad-hoc sensor networks and pervasive computing. While working at Philips Research, I have also developed interests in Digital Rights Management (DRM).

1.1 Autonomic Management of Ubiquitous Systems

In the healthcare industry, we witnessed the development of numerous sensors for monitoring physiological parameters including pulse, heart-rate, oxygen saturation, as well as behavioral parameters such as posture and gait. On the other hand, there is considerable research on the design of new body sensors and measurement techniques and of actuator devices such as drug pumps, bio-electrical and bio-mechanical devices. These sensors and actuators are usually wearable and often implantable. They can form body-area networks by communicating wirelessly among themselves and can interact with wearable processing units such as PDAs, mobile phones, and the fixed network infrastructure in the environment. Such body area networks provide a good basis for continuous monitoring of a patient's medical condition anytime and anywhere. For example, medical sensors (developed for the body sensor nodes, e.g., Telos, MicaZ and BSN motes platform) can be programmed to monitor physiological characteristics of patients in post-operative care, patients with episodic manifestations (such as cardiac arrhythmia) and chronically ill patients (e.g., those suffering from diabetes mellitus) as they go about their daily activity. The benefits to the patients include early release from hospital and constant monitoring of their clinical condition as well as automated alerts and assistance from healthcare personnel when needed. However, whilst the miniaturization and the proliferation of various sensors and devices for e-Health implies that they can effectively disappear, the management and configuration of multiple sensors and software components without explicit user or administrator intervention poses many challenges.

While working as a postdoctoral researcher at Imperial College, I collaborated closely with several researchers from the Biomedical Engineering department to work on two research projects to continuously monitor patients with chronic diseases. They were responsible for designing new biomedical sensors, while we from the Department of Computing looked at the system management and security aspects of body sensor networks. In particular, we advocate that (1) Devices and sensors should dynamically discover each other and then determine whether interaction between them is required. (2) Their configuration may need to change depending on the user's context, e.g., location, current activity and medical history. (3) physiological parameters for clinical observation

and monitoring, such as heart rate thresholds need to be configured and customized. This in turn leads to the need for adaptation according to the changes in context and clinical conditions over time. Thus, we implemented a policy-based architecture that supports autonomic management for self-configuring and self-managing such systems, using the concept of a *Self-Managed Cell* (SMC). An SMC consists of autonomous set of hardware and software components that represent an administrative domain such as a body area network of physiological sensors and controllers. SMCs are able to function autonomously and adapt automatically to the user's current activity, environment, communication capability as well as interactions with other SMCs. An SMC implements a local feedback control loop where changes of state in the managed objects and resources trigger adaptation that in turn affects the state of the system.

An SMC manages a set of heterogeneous components (i.e., managed resources) such as those in a body-sensor network, a room or even a large-scale distributed application. Resource adapters are instantiated to provide a unified view for interaction with managed resources as they may use different interfaces for communication and invocation of management actions. SMCs may also comprise additional services for detecting context changes, monitoring of component behavior and providing security (e.g., authentication and intrusion detection). However, a number of services constitute the SMC's core functionality and must be present in every SMC instantiation. These include an event bus, a policy service, and a discovery service.

The integration of the services and components occurs through a common publish subscribe event bus. This has the advantage of decoupling the services, since an event publisher does not need to have prior knowledge of the recipients when sending the message, and permits the addition of new services to the SMC without disrupting the behavior of existing ones. The event bus can be used for both management and application data such as alarms indicating that threshold have been exceeded. The policy service implements a local feedback control loop to achieve the SMC's adaptation requirements in order to realize self-management. It caters for two types of policies: obligation policies (event-condition-action rules) define the configuration actions that must be performed in response to events, while authorization policies specify what actions are permitted on which resources and services. Policies can be added, removed, enabled and disabled to change the behavior of the SMC components without code modifications and interrupting the SMC's functioning. The discovery service is used to discover components which are capable of becoming members of the SMC in the vicinity. It is also responsible for managing group membership as it is necessary to cater for transient failures, which are common in wireless communications, and permanent departure.

Although we focused on e-Health applications in this research, the architecture is meant to be applicable to a wide range of pervasive and other applications such as smart spaces, vehicular networks or unmanned autonomous vehicles, etc. We consider the SMC as an architectural pattern that can be tailored on instantiation and that can be applied at different levels of scale from body-area networks to large scale distributed systems.

1.2 Security for Wireless Ad-hoc Networks and Ubiquitous Computing

The recent advancements in mobile technology and wireless transmission have led to the increased development of ad-hoc networks that can be formed at anytime and anywhere. Ad-hoc networks are frequently formed to enable collaborations and interactions between mobile devices or mobile vehicles in order to share resources, information and services. In the future, devices maybe embedded into electrical appliances, in the infrastructure of buildings, vehicles and in everyday artifacts.

They will have the capability to interact with each other by relying on each other to provide the required resources and services towards the purpose of the collaboration. This leads to the portrait of the 21st century computing by Mark Weiser, i.e., ubiquitous computing. Ad-hoc networks can be seen as one of the enabling technologies to realize the vision of ubiquitous computing.

My Ph.D. thesis proposed a security framework to enable the formation, evolution and management of wireless ad-hoc networks consisting of autonomous mobile devices. It integrated different strands of research including a distributed role-based access control model and light-weight interaction protocols for the secure establishment and evolution of ad-hoc networks. A policy-based model was used to control the behavior of participants in the network both in terms of authorizations and of adaption to changes of context.

My thesis addressed some fundamental challenges of ad-hoc networks: (1) How to form an ad-hoc network, (2) How to determine who can participate and join the collaboration, (3) what resources and services are needed to form the network, and (4) who are authorized to access the resources and services of the network. A community specification called *doctrine* was introduced to specify membership admission policies, authorization policies to control access to the services, as well as obligation policies to manage and adapt to context changes in the network. Policies have been used in many application areas for managing distributed systems, active networks and pervasive computing. A policy-based approach was proposed because it is flexible, scalable and permits adaptation to context changes in security requirements by dynamically loading and removing policies from the system without interrupting its functioning.

The thesis also proposed that participants must be authenticated before they can be granted admission. This involves the verification of their credentials in order to ensure that they satisfy the membership admission policies. However, wireless ad-hoc networks typically have very distinct characteristics from the wired networks in that they cannot rely on the availability of a fixed network infrastructure. Therefore, a continuous connection to the Internet cannot be guaranteed at all times and this implies that Certification Authorities (CAs) may not be reachable, the public-keys of CA and the revocation status of a participant's certificate cannot be checked. Consequently, participants can only rely on their own knowledge as well as additional security information held by other peers in order to verify each other's credentials. We introduced the *trust-based credential verification scheme* to provide a flexible approach that allows participants to exchange information regarding credentials that they have verified and the validity of credentials that they know of, i.e., through the use of *a priori* knowledge and *out of band* verification.

Membership management in ad-hoc networks is required in order to ensure that only authorized participants are permitted to join the network and to access the resources or services. Existing approaches are either too expensive due to the computationally expensive threshold cryptography in addition to the laborious process of setting up CAs or do not maintain the membership at all. The proposed framework relies on a coordinator to be selected among the participants in order to enforce policies. Instead of using a certification scheme, the coordinator broadcasts the membership list to all participants at regular interval using TESLA authenticated broadcast in order to maintain weak consistency of the membership in the network. In order to address the issue of single point of failure, an efficient replacement scheme to replace the coordinator had been proposed. Our approach significantly reduces the redundancy in the verification of credentials whenever a new participant requests to join the network. In addition, the membership list serves as both the membership certificate and revocation list (CRL), thus eliminating the need to periodically broadcasting a CRL.

Security and management protocols were also designed in order to realize the formation of ad-hoc

networks and to ensure the authenticity and integrity of the information exchanged between participants. The protocols include disseminating the policies, discovering the devices (participants), disseminating security related information, as well as handling join requests from participants and disconnection of existing members.

We have also studied various access control models, in particular role-based access control (RBAC) has proven to be effective and is able to reduce administrative costs. We have successfully incorporated role based management into ad-hoc networks in order to regulate the behavior of participants in terms of authorizations and of adaptation to context changes such as changes in the network topology, locations and occurrence of security violations.

2 Future Research Plan

In this e-Health era, with the invention of numerous medical sensors and devices that are capable of gathering medical measurements for the purpose of monitoring many physiological parameters of patients and the elderly, the vision of Ambient Assisted Living (AAL) can soon be achieved. AAL enables the patients and the elderly to live normally in their own home, at the same time allows for their medical conditions to be monitored continuously. However, in order to achieve this goal, there is a strong need to form a close collaboration between the bio-medical/electrical engineers to invent new types of medical sensors, the human-computer interaction and usability researchers to design a natural and user-friendly way of using the medical devices, and the computer scientists to construct a scalable software architecture to arrive at a seamless integration of such monitoring system into the environment. On the other hand, security for such healthcare applications cannot be neglected. At the later stage of my postdoctoral work, we devised a protocol to enable secure discovery and sensor-to-patient's controller association (pairing) while bootstrapping a body sensor network. This is essential as we must ensure that the medical sensors are associated with the right patient. As the communication between sensors occurs through wireless medium, (e.g., ZigBee), the medical measurements, commands issued by the actuators in respond to the changes in the patient's physiological parameters, and event notifications to cater for network configuration changes are susceptible to numerous attacks such as passive eavesdropping, message modification, impersonation and masquerading. The cryptographic libraries specially tailored for resource constrained sensor networks have been invented such as *TinySec* and *TinyECC*, they can be exploited further to serve as the foundation for innovating efficient and lightweight cryptographic protocols as well as key management schemes to address numerous security threats of e-Health monitoring.

Data gathered from the various sensors and medical devices at home is not restricted to in-home monitoring only, but they could also be relayed and stored in the medical data centers in the cloud for further diagnosis and analysis. These data are extremely useful for monitoring the lifestyle of a healthy individual, a community or a patient (e.g., patients with high blood pressure) whether he fulfills the obligation to undertake a specified amount of exercises a day (e.g., monitoring of daily activities can be done using *Philips's DirectLife*, *Wii fit*, and *Adidas's MiCoach*, etc). Consequently, some level of coaching can then be given to the users or the patients based on their performance and their fitness targets. Additionally, the medical data gathered could be used for discovering the diseases pattern, and possibly used for predicting a pandemic outbreak through the use of data mining and machine learning techniques. The privacy of users/patients must be preserved for such systems in order to prevent unauthorized parties from learning about the user's behaviors such as their daily activities, and diseases that a person might have contracted. Many privacy-preserving techniques could be applied here, e.g., using the homomorphic encryption

scheme, statistical inferencing and zero-knowledge proof, one could perform some operations such as searching on the encrypted data in order to execute some form of inferencing. This is especially challenging given that the data mining and inferencing operations on the encrypted data need to be time-efficient. On the other hand, from the device authentication perspective, it is important to ensure that the device is strongly coupled with the entity being monitored, possibly through biometric identification such as fingerprints and gait identification. Otherwise, impersonation attack could be launched by attaching the sensors/devices on a healthy person, thus providing “healthy” measurements to the monitoring system.

While working at Philips Research on Digital Rights Management (DRM) for multimedia content management, the notion of applying DRM to healthcare applications triggers some research interests. Ultimately, with the mandate of the use of Electronic Health Record (EHR) in the US starting in 2014, it implies that the EHR can be replicated and disseminated easily, patient’s health record is no longer a matter of keeping the paper record in the cabinet locked and hence no one has access to it. In addition, the health data can be stored in many distributed services (such as Google Health and Microsoft HealthVault) in different form, and defining access policies for this sensitive data is not trivial. When the enforcement of access control is shifted from the server (data center hosting the EHR) to the client (entity requesting access to the EHR), e.g., when a patient visits a dietician in a private clinic, the dietician can obtain the encrypted EHR from the Google Health Service or any other relevant source; the permission to access the EHR can then be requested on the fly by obtaining a consent (or “license” in the DRM context) from the patient in order to decrypt the EHR. Contrary to traditional access control model enforced at the server side that requires access control policies to be defined *a priori*, the DRM approach provides the flexibility to the patient to grant consent to whichever parties that he believes “should” have access to the EHR for obvious reason. However, this comes with a price in that the need for a tamper-resistant hardware on the client device then becomes a strong requirement in order to ensure that the decrypted EHR content is not leaked to the third party.

Apart from the healthcare domain, I believe that the era of Ubiquitous Computing has provided a great opportunity to learn more about the user’s behavior. A trace of the user’s behavior can be captured through sensory input and user actions in order to gather information about how a user behaves and reacts in different circumstances, locations, and encounters with other entities as well as changes in user context. Typically, policies can be used to govern access to user’s information, but policy specification is a difficult task and prone to error. Therefore it would be nice if we can design a system that can learn new policies through observation of the user’s behavior in order to predict future behavior and to facilitate automated user actions with minimal user intervention. However, the learned policies must be refined regularly and adapted because the user’s behavior may change over time. Although policies are used to govern the access to the user’s information, from time to time, the user can explicitly overwrite the policies to allow for exceptions such as whenever there are new constraints that have not been catered for in the past. An audit log maybe used to capture these events, thus enabling the policies to be refined in order to predict the user’s behavior more accurately. Data mining and artificial intelligent techniques can be applied to support incremental learning of policies, as well as to resolve conflicts to previously learned policies. In most cases, the user wouldn’t know how she will react in certain events/requests *a priori*, therefore dynamically creating policies based on the user actions can autonomously guide the user’s future behavior.

Lastly, this document presented some exemplars of research areas and ideas that I plan to engage in in the future. With my previous research experience both in academia and industry, I believe I can make substantial contribution to the research in these areas.