Mapping the Physical World into the Virtual World: A Com2monSense Approach

Richard Gold, Vlasios Tsiatsis, Jan Höller, and Vincent Huang

Ericsson Research

Abstract. Networks of the future will contain a vast array of sensors and actuators which users will wish to interact with. Orchestrating this task effectively and efficiently will be one of the major challenges facing the deployment of these future networks. We propose a three-tier architecture which provides a middleware to resolve high-level service requests from applications to low-level sensor or actuator actions. Our architecture is designed around the core principle of horizontalization, that is, breaking apart vertically integrated silos into their component parts, thus allowing flexible recombination and reuse of data & services.

1 Introduction

Research in SANIs (Sensor Actuator Network Islands), i.e., a group of sensors with networking capabilities communicating with a gateway, has typically been concerned with issues such as: radio design, operating systems, network protocols and energy efficiency of highly specialized (e.g. enterprise, environmental) and small to medium scale sensor network islands [5]. The focus of this research is more on the deployment issues (e.g. coverage, energy efficiency [4] etc.) and less on the purpose of the deployment of sensor networks which is to provide an information interface from/to the physical world. This research has been extremely useful, but its narrow focus necessarily requires that the larger picture of how sensor networks are used has not been equally as well addressed. Specifically, the actual applications, use cases and scenarios in which sensor networks will be used will be a major influence in the design of sensor networks. Additionally, the networks of the future will be deeply heterogeneous (Section Assumptions) which will affect the design of any architecture for usage in such networks. In this paper we present an alternative outlook which focuses on how sensor networks will actually be used. Our research agenda is not focused on traditional areas of sensor networking research as outlined above. We are, however, primarily interested in an infrastructure that ties together the multitude of sensor and actuator deployments and exposes the collective set of services to be used by applications. The main contribution of this paper is the design, analysis and prototype of such an infrastructure constructed around the design principle of horizontalization by which we mean a horizontal market where the services offered by the sensor networks can be combined and reused by many different apapplications. This approach can also be thought of as a middleware for sensor networks. Currently sensor networks are typically deployed in vertically integrated silos which are self-contained. This means that if new functionality is required, then an entirely new silo must be deployed. Silos also represent information separation which

C. Hesselman and C. Giannelli (Eds.): Mobilware 2009 Workshops, LNICST 12, pp. 142–146, 2009.

is in contrast with the information aggregation and and fusion that is desirable from sensor actuator network islands. We see this as inefficient and one of the main goals of our architecture is to allow the reuse and re-combination of the data and services of sensor networks without requiring the deployment of new silos. We have already begun preliminary implementation work based on Middleware and Semantic Web technologies. We believe that our architecture is applicable to a wide range of scenarios and we describe and analyze a carefully selected example.

2 Com2monSense

The Com2monSense architecture is based on the middleware approach outlined above and was partially presented in [1,3,2]. The Middleware layer is the central point of the system, providing a natural point for the sensor networks and the applications to connect to. Additionally, it also provides an opportunity to bring 3rd party services into the interactions between the sensor networks and the applications, potentially providing services such as maps and tag resolvers for RFID. One of the direct consequences of this approach is the breaking up of the vertical silos of the Client/Server architecture into its separate components. This means that the data and services of the sensor networks are now re-combinable. Another consequence is enabling the creation of new functionality without deploying an entirely new silo.

One possible realization of the architecture described above is shown in Figure 1.

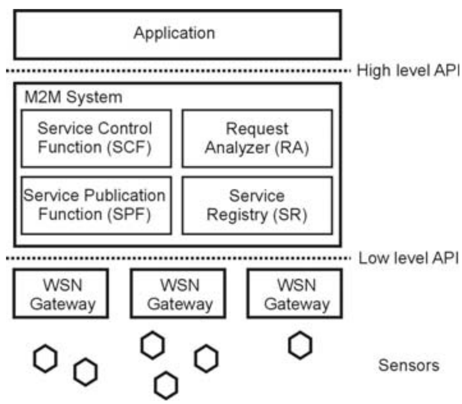

Fig. 1. High-level System Overview

The main components of the system are as follows:

Application: The application resides on the client device and is responsible for representing the clients' preferences and performing the user's desired actions. The application sends high-level service requests to the M2M system.

The M2M System: This is comprised of the following parts:

- **Service Control Function (SCF):** This is the main multiplexing & demultiplexing point of the M2M System. High-level service requests come in to the SCF from the

application and low-level function call responses from the sensor network gateways come in to the SCF. The SCF then needs to multiplex the service requests to the appropriate part of the M2M System and demultiplex the function call responses to the correct part of the M2M System.
- **Request Analyzer (RA):** The RA is the key part of the M2M System. It is responsible for breaking down the high-level service requests into low-level function calls which are to be sent to the gateways.
- **Service Registry (SR):** The SR keeps track of which sensor networks are connected to the M2M System, which capabilities they have and how to reach them (e.g., which network technology is used to communicate with them).
- **Service Publication Function SPF:** This is a module for future use and we do not consider it further in this paper.

Gateways: The gateways manage the attachment and subsequent detachment of sensors and the general network configuration. The sensors will report their capabilities to the gateway which will then report the set of combined capabilities of its attached sensor networks to the M2M System as a set of services.

Sensors: The sensors register at the gateway automatically when they are powered on and provide a list of their capabilities. This should be performed with some standardized auto-configuration protocol. They then either periodically report information to the gateway or respond to explicit requests forwarded to them via the gateway according to how they are configured.

3 Implementation

In order to realize the Com2monSense architecture, a variety of technologies need to be employed. Since we are aim to support multiple applications and sensor networks, it is critical that we have a standardized set of interfaces:

- Interface from the application to the middleware (for example, SOAP, SIP or XMPP)
- Interface between 3rd party providers (for example, AJAX)
- Interface from the middleware to the sensor networks (For example, SIP or XMPP)

Once these are in place, the next step is the creation of the middleware itself. The primary goal of the middleware is the translation of high-level user queries to low level SANI calls by combining information from multiple sources to present a coherent view to the user. We believe this to be the part of the architecture that requires the most work. Currently we are looking at Semantic Web technologies to assist us in this area. A goal of the Semantic Web to make information machine-processable is also useful for helping us resolve high-level service requests to low-level sensor network function calls. Finally, the goal of the Semantic Web to integrate data from heterogeneous sources makes it a natural fit for our ambition to combine data from multiple sensor networks for consumption by multiple applications.

We use the ontology capabilities of the Semantic Web to represent the various aspects of the physical world. For example, sensors in different rooms, cars, streets, buildings, etc. can be represented in an ontology with explicit links between them. The motivation for having these links between the representation of sensors is that we wish to use *inference* as method of working out which sensors can be used to satisfy which service requests, i.e., working out how to break down high-level service requests to low-level function calls.

4 Scenario: Smart Open Spaces

We envisage the instrumentation of public, open spaces with various types of sensor to enable a rich set of services for the public. This will take the form of mapping the physical space into the Internet by allowing the public to interact with the physical environment through the Internet via the use of mobile devices. Interaction with the local environment can take place both locally and remotely. It will be possible for the virtual representation of the physical environment to be queried, annotated and remodelled by each individual user. The user will be able to discover the physical state of the environment (weather conditions, etc.) as well as access services such as gaming, tourism and community services. When a member of the public visits the smart open space, they may wish to know in advance the current status of the open space. For example, how many people are already there? What is the weather like? If the user then proceeds to visit the open space, they may wish to view annotations left by other people regarding Points-of- Interest and nearby services such as restaurants or museums.

The smart open space provides a horizontalized system where multiple service providers can attach to the open space and advertise their services and information models which users can then selectively choose or enable. This scenario represents an enhancement of open spaces for the general public. It also provides a neutral platform for companies to provide services and/or information. It will make it easier and quicker for the general public to discover better information about their physical environment and also use a rich array of services to interact with the environment. Due to the face that information is being collected about people's movements in public spaces, privacy is a major concern.

5 Conclusions

We have presented a three tier architecture for the processing and presentation of sensor data called Com2monSense. Our architecture is based around the core design principle of horizontalization, that is, breaking apart vertically integrated silos into their component parts thus allow recombination and reuse of the data and services of sensor networks. For future work we plan to continue our prototyping efforts to create a robust implementation of our architecture and examples applications conforming to the presented scenarios. Also, we plan to analyze in-depth the issues of privacy and security which we believe are essential to widespread uptake of such an approach.

References

1. Krco, S., Johansson, M., Tsiatsis, V.: A CommonSense Approach to Real-world Global Sensing. In: Proceedings of the SenseID: Convergence of RFID and Wireless Sensor Networks and their Applications workshop, ACM SenSys 2007, Sydney, Australia (November 2007)

2. Huang, V., Johansson, M.: Usage of semantic web technologies in a future M2M communication system. In: Proceedings of the 1st European Semantic Technology Conference, Vienna, Austria, May 31- June 1 (2007)
3. Huang, V., Javed, M.K.: Semantic Sensor Information Description and Processing. In: Proceedings of the Second International Conference on Sensor Technologies and Applications, SENSORCOMM 2008, Cap Esterel, France, August 25-31 (2008)
4. Hadim, S., Nader, M.: Middleware Challenges and Approaches for Wireless Sensor Networks. IEEE Distributed Systems Online (2006)
5. Römer, K., Friedemann, M.: The Design Space of Wireless Sensor Networks. IEEE Wireless Communications (December 2004)