

Application-Oriented Fusion and Aggregation of Sensor Data

Gee Fung Sit¹, Chenbin Shen¹, Holger Storf², and Cristian Hofmann¹

¹ Fraunhofer Institute for Computer Graphics Research,
64283 Darmstadt, Germany

{gee.fung.sit, chenbin.shen, cristian.hofmann}@igd.fraunhofer.de

² Fraunhofer Institute for Experimental Software Engineering,
67663 Kaiserslautern, Germany

holger.storf@iese.fraunhofer.de

Abstract. A glance at the today’s research and industry community shows that AAL installations are normally offered as “complete solutions”, often including overlapping of almost equal or homogeneous sensors. Thus, redundant sensors are integrated in one single space when purchasing different AAL solutions, leading to an increase of acquisition costs and higher data volume. In order to counteract this problem, we present a method for application-oriented fusion and aggregation of sensor data. Here, the main contribution is a reference model and a semi-automatic approach for the determination of applicability of sensors to predefined AAL applications.

1 Introduction

In specific areas such as home automation, activity recognition, or smart metering, a multitude of sensor and actuator technologies are employed which have different properties, e.g., wireless vs. wired communication, microsystems and terminals, etc. Here, methods for sensor fusion and aggregation have a supporting effect: the combination or junction of sensor data leads to better quality of gathered information in comparison with the consideration of individual devices. In this scope, improvement of information refers to a more stable behavior with perturbations and an increased clearness of statement by means of enhancement of the measurement range and the resolution of measured data [1]. Furthermore, additional information, like characteristic activities of daily living (ADL) for long-term behavior monitoring, can be gathered from the combination of sensor data, which cannot be captured from the single data streams [2][4]. So, main directions of sensor fusion focus on activity recognition, context recognition, or personal identification.

A consideration of the research community shows that, from a methodological point of view, especially the area of sensor technology for activity and context recognition is not sufficiently understood and supported. Current publications on AAL applications indicate that sensor data is often integrated in an application or system in a proprietary way (tailored to specific requirements) or, if basic

interfaces are defined, no detailed specification of sensor types and data formats are provided. Common users and deployers of respective systems, for instance programmers, system integrators or engineers still rely on ad hoc defined interfaces and methods of data processing. Thus, AAL solutions or installations are normally offered as complete packages consisting of a collection of sensors (and actors) along with software supplying complex behavior. As these are tailored for specific needs it stands to reason that multiple systems may be required within an AAL space. In essence, each particular need, i.e. each application, requires the acquisition of an entire AAL package. Since the physical devices are typically fairly basic and straightforward in their behavior, this may lead to redundancy whenever multiple solutions are employed. This duplication leads to higher costs (both initial and running) caused by redundant purchase, as well as an increased maintenance. Additionally, data overhead is increased, which may result in excessive wiring or, if wireless technology is used, lead to bandwidth problems and increased interference.

In this work, we aim at reducing the overlap in hardware devices (sensors and actuators) using a semi-automatic approach. The main goal can be described as follows: Given an AAL platform and an application selected by the user, the system is to provide feedback on whether the application can be realized with the devices at hand. If no matches are identified, the system is to signalize which devices need to be bought additionally. The groundwork for solving this problem was laid with the conception of AAL interoperability platforms such as Continua, OASIS and universAAL. Basically, they allow individual sensors and actuators to be replaced by semantically equivalent devices to ensure that the information and services provided by sensors and actuators from the physical devices are in a sense decoupled from the physical devices. Here, we show that, given a suitable data model, it is possible to extend this approach to higher abstraction levels.

The remainder of this paper is organized as followed. First we present related work with respect to the identification of sensors and actuators based on a predefined application. Second, we present the conceptual model for application-oriented fusion and aggregation of sensor data, generally consisting of a specific data model and reference model as well as a workflow-based description of a three-step semi-automatic approach. Third, we show how this approach can be employed in the scope of activity recognition. Finally, we recap the results and show venues for further research.

2 Related Work

There are several approaches to model and implement a universal AAL platform to ensure interoperability between different sets of hard- and software components. Here, we particularly need to consider approaches which deal with the junction of sensor data in a common context, and allow the retrieval of sensor types that are applicable in specific scenarios. Regarding this issue, only a few work concerning sensor fusion and aggregation can be identified.

Within the AAL research community, the OASIS project deserved mention for introducing the concept called hyper-ontology which is used in conjunction with ontology alignment to map applications to their domains [5].

Outside of AAL the Open Geospatial Consortium's (OGC) Sensor Web Enablement standard offers sensor fusion with an emphasis on geospatial data.[6] It relies on the OpenGIS Sensor Model Language (SensorML) which describes high level processes as partitioned process chains. The standard has an extensive system for describing each constituent semantically, particularly its input, output, and the underlying physical laws or algorithm that governs its inner workings [7].

3 A Conceptual Model for Application-Oriented Fusion and Aggregation of Sensor Data

In this section, we introduce a conceptual solution for application-oriented fusion and aggregation of sensor data. First, we present an overview of the approach and substantiate its feasibility in short. Then we discuss the data and reference model for a three-step semi-automatic approach.

3.1 Overview

The concept presented herein relies on the existing AAL platform universAAL [10]. This choice is meant to be representative for other state-of-the-art AAL systems: First, it comprises a data model based on the RDF standard [11] for the purpose of information exchange. Second, we expect resources in the AAL space - in both the physical and virtual realms - to be modeled by means of an ontology, e.g., using OWL. Here, OWL-based messages ("context events") describing the current state of (a part) of the AAL space ("context") can be provided as an RDF statement.

In the user-centric view, an application is a black box that provides a certain set of functionality. In contrast, we concentrate on the context the application operates in and the services it needs to call. We call this the *signature* of the application.

Observe that contexts can be atomic, given at sensor level, or derived via sensor fusion and algorithmic reasoning. A similar statement can be made for services. In the atomic case, the software would be dependent on the given hardware, in the second case however the application would be agnostic to the physical devices. In its simplest form, this corresponds to a common approach in IT systems to achieve interoperability: software wrappers and drivers abstract over a homogenous set of software and hardware respectively. In general, we need to abstract a heterogeneous set of a large number of possible setups with many-to-many relationships between the resources and their abstraction. Figure 1 outlines an abstract example showing one layer of abstraction over the device context. In reverse direction, given an arbitrary context, we can check if it can be broken down into simpler ones. We call this action a *decomposition*.

Figure 2 illustrates this concept. The main idea is to describe the application in the most general terms and consider possible decompositions. This means that reasoning and sensor fusion components are pulled from the applications to gain a greater independence from low-level details. Similarly, a service can invoke another service. Accordingly, there are specific dependencies between different services.

The approach presented herein is semi-automatic in the sense that human input is used to resolve ambiguities that may arise in certain situations. To give an example, Let us assume that we want to monitor a person, raising an alarm when he or she falls. To realize this functionality, we can either fit the client’s body or the floors in his home with sensors. A priori, both options are equally valid. This represents a real choice for the user.

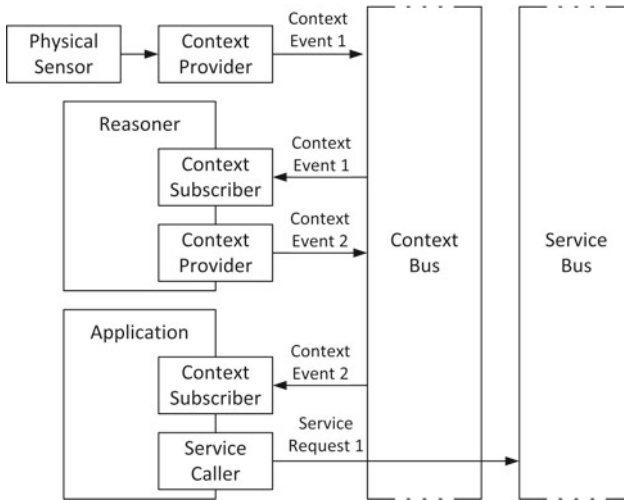


Fig. 1. Data Flow in a universAAL Application

3.2 Feasibility

As discussed in the introduction, the current paradigm is having a complete system of soft- and hardware for every application. The user determines which applications he or she requires and, accordingly, acquires packages that provide the specific functionalities. Implicitly, we get a mapping from the sensors, actuators, and governing software to the provided features. This information can also explicitly be provided in external repositories by manufacturers and developers, various stakeholders, such as health care staff or patient groups, and third parties, such as research groups.¹

¹ cf. OASIS hyper-ontology concept in [5].

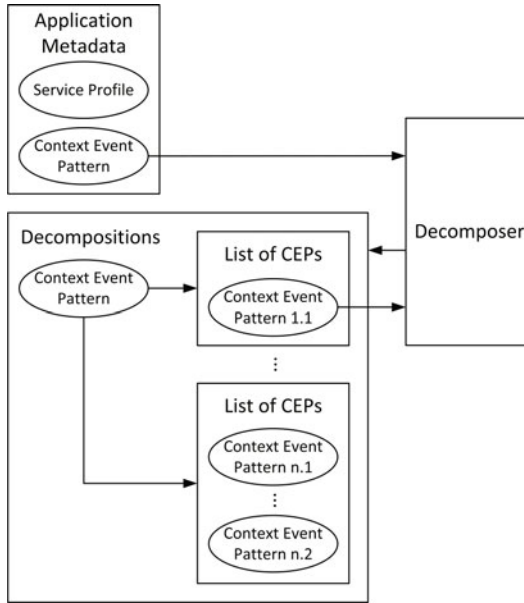


Fig. 2. Data Flow in Decomposition

3.3 Data Model

Since this project builds upon the universAAL platform, we inherit its data model. As already mentioned in Section 3.1, universAAL is meant to be representative for state-of-the-art AAL platforms. In essence, we operate on an OWL ontology and capitalize on the RDF model. These standards ensure that the concept presented in this paper can be ported to support alternative platforms. First, we recap the most important facts about the data model. Then we present the metadata model for applications built upon the models already given.

3.3.1 Context and Services

UniversAAL uses two different buses for communication between its components. The context bus is event-based and used to provide data. The service bus is request-based and used to access functionality. Recall that context events are given as RDF statements. They are specified via *context event patterns*, essentially lists of OWL restrictions over its properties (i.e. subject, predicate, and object). UniversAAL has three different types of resources that can provide a context: gauges, controllers, and reasoners. A gauge is a wrapper for a physical sensor, an abstraction layer that transforms its raw output into a context event. Similarly controllers abstract over actuators and create an event based on the change it facilitated. Finally reasoners infer context based on other context events. Service metadata is similarly encapsulated as service profiles (cf. [12]). Service profiles characterize a service with two basic types of information: “What

is the service for” and “how it can be utilized” [12]. Here, we concentrate on sensors and accordingly on context event patterns. Whenever possible we provide the corresponding procedure for handling services.

3.3.2 Ontologies

Every object in universAAL, both physical and virtual, is modeled as an RDF resource or a subclass thereof. Each resource within an AAL space has a unique identifier. To capture the complex relations and interactions between the components, ontologies are employed. The first trait ensures that both context event patterns and service profiles are well-defined, if used properly. The second one allows us to qualify arbitrarily intricate sets of resources via restrictions.

3.3.3 Application Metadata

We characterize an application by its signature as defined in Section 3.1. That is, it is stored as a list of context event patterns and service profiles. By using the same data structure as the AAL platform, we ensure that we can easily check against the state of the AAL space and the low level services it provides. Given that ultimately we rely on the common standards RDF and OWL, it is possible in principle to translate between different AAL platforms that fulfill the minimal set of conditions given in Section 3.1.

3.3.4 Decompositions and Dependencies

The concept of decompositions as introduced in Section 3.1 entails a one-to-many relationship between a higher-level context and a list of lower-level contexts. In the case of contexts, this description closely related to the concept of the reasoner which takes lower-level concepts to compute a high-level one; both follow from the underlying concept of abstraction. Note, however, that the abstraction level is not a formally defined property. While a tree-like structure arises naturally when we consider a single context by itself, as a whole we can only characterize the relations as a directed graph. Therefore, care must be taken to avoid or detect cycles. The same is the case for resolving dependencies between services. These arise when a service invokes other services.

3.4 Reference Model

The concept presented here can be divided into three different parts. In the first one we discuss the required data stores. Then we present architectural components which set up the reference model. Finally, we explain a three-step semi-automatic process, examining a respective workflow in detail.

3.4.1 Data Stores

As mentioned in Section 3.3.3, we need to store application signatures. We call the required data store the Application Database (see Figure 3c). Here we also store a human readable string for identification. The Application Database can

be realized as a local database. Additionally, external repositories are required for decompositions of contexts and dependencies between services, and ontologies to retrieve mappings from context event patterns and service profiles to physical devices and accompanying software. We presume the existence of a multitude of repositories, reflecting diversity in devices and the number of their manufacturers.

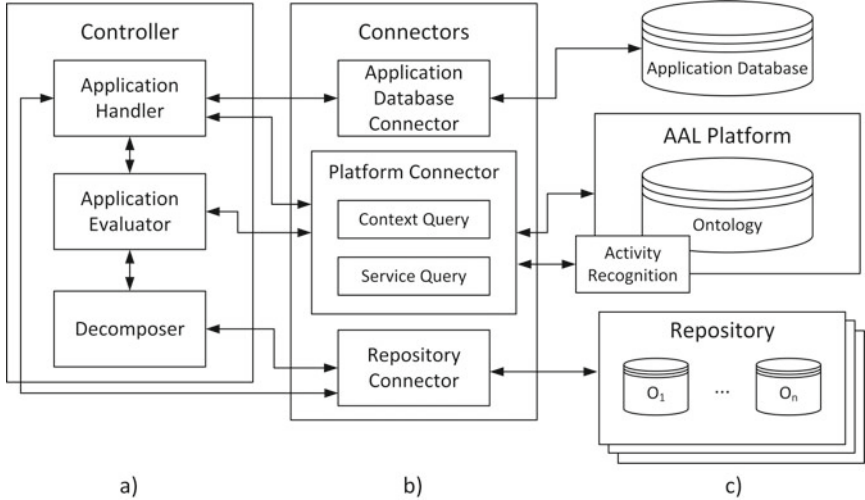


Fig. 3. Reference Model

3.4.2 Components

At the lowest level there are so called Connector components, namely Platform Connector, Application Database Connector, and Repository Connector (see Figure 3b). The Platform Connector comprises interfaces that enable the communication with the AAL platform we operate on. First and, foremost these components are responsible for making inquiries about the build-up and the state of the AAL space. Additionally, we have connectors to the external repositories and to an application database that holds the applications the user is interested, both introduced in the previous section. On the next level we have the Controller layer, which consists of Application Handler, Application Evaluator, and Decomposer. These components are used to handle the input from the user and the output of the Connector components (see Figure 3a).

In this section, we described the basic components included in the reference model in abstract terms. The specific functions of these components are elucidated in the following section by means of the workflow associated with a three-step semi-automatic process.

3.4.3 Workflow

On the basis of a specific workflow we demonstrate how application-oriented sensor data fusion and aggregation is realized by the presented reference model in a three-step semi-automatic process. Given an application, the first step is to check for the existence of the contexts and services that are required for it to run. If a resource is missing external repositories are consulted. In the final step the user receives a list of options of devices and software that would allow him to employ the application in question. In this section we discuss the individual roles of the components introduced above in order to describe how to accomplish the tasks explained below. Figure 4 shows the corresponding data flow.

From an application-centric point of view the following tasks have to be considered: fetching an application from the Application Database (marked with dashed arrows in Figure 4), creating new applications (marked with white arrows in Figure 4), and evaluating applications to determine whether they can be employed within the current AAL set-up (marked with black arrows in Figure 4).

The Controller layer, as mentioned above, handles all the communication between the user and the presented system, hiding the low-level data exchange with the AAL platform and the data stores. The Application Handler is the key component for user interaction. First, it loads a list of the existing applications from the Application Database using the according connector component. Second, it can be used to create new applications by combining existing resources in the AAL space. To accomplish this, the Application Handler employs the Platform Connector and the Repository Connector to execute queries over the existing local resources and external ontologies respectively (cf. Section 3.4.1). The user is provided with a list of possible contexts and services from which a subset can be chosen and aggregated. The new application can now be stored in the Application Database.

However, before it can be deployed, we have to ensure that the required resources are available. By definition these are fully specified by the Application signature (see Section 3.3.3). To check the signature against the AAL setup, the application metadata is forwarded to the Application Evaluator. This component performs queries for the context event patterns and service profiles using the appropriate sub-components of the Platform Connector. If every signature entry is well matched with a corresponding resource in the AAL space, a positive result is returned to the Application Handler.

If the result is negative, i.e., not every signature fits to a resource, the unmatched resources are passed to the Decomposer component in order to start the decomposing process (see Figure 2). As mentioned in Section 3.3.4, higher-level context events are the result of sensor fusion or aggregation while services form dependencies. In the latter case, missing services are searched for in external repositories via the Repository Connector using the associated service profile as parameter. Depending on the success of the query, the Decomposer either returns a list of possible candidates to the Application Evaluator or reports the negative result. In the case of contexts, a more differentiated approach is

required. First let us consider the simplest instances. A context may be entirely missing from the repositories if it is malformed or suitable soft- and hardware has not yet been developed or released to the public. Here, we simply report a negative result to the Application Evaluator. Another possibility would be having an atomic context. In this case, the context event pattern can be matched to devices which provide the associated context events. This list is simply returned to the Application Evaluator.

Recall however, that in general there is a one-to-many relationship between a context and possible decompositions. The list of decompositions can be presented to the Application Evaluator. In this case, each entry in each decomposition list could in principle again necessitate a decomposition process. Since the underlying data structure is recursive, in theory we have to rely on the repository owners to make sure the queries terminate properly. In practice, a recursion depth limit should be implemented along with cycle detection.

After this recursive fusion and fission process, the Application Handler should get the message about which devices can or must be integrated in order to realize the given application. Thus the Application Handler is able to give user-feedback about the devices that have to be purchased.

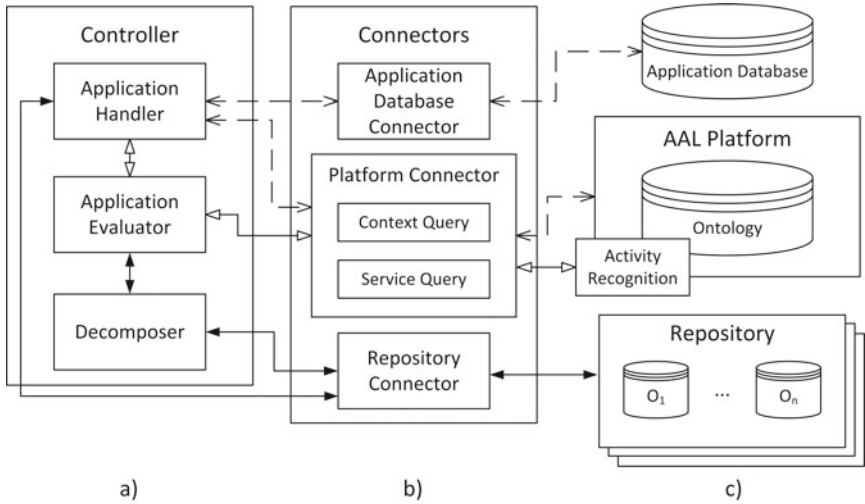


Fig. 4. Reference Model with Data Flow

4 Fusion and Aggregation for Activity Recognition

In the previous Section, we presented a model which enables the determination of adequate sensors based on a predefined AAL application. Given the assignment of sensors to applications, fusion and aggregation can be performed by means of specific reasoning components.

As already mentioned in Section 1, sensor (context) events can be joined to composed events through fusion and aggregation and, further on, to higher-level activities. With respect to the presented reference model, an example for a connected component (see Figure 3c) is the Event-based Activity Recognition and Reaction System (EARS) which has been initially developed in the context of the EMERGE project [8].

The major goal of this component is to identify specific patterns at the continuous event stream. This will be done on-the-fly, so that it is possible to trigger an adequate reaction if necessary. Typical event patterns from interest are the so called (instrumental) Activities of Daily Living (ADL) [3] for behavior monitoring on a long-term view. The challenge is that the patterns of ADLs can be very different in their structure like sequential or fuzzy or a mixture. Additionally the activities can be individually performed and may overlap.

Another requirement for the EARS-framework was that learning of the activities should be avoided because of its obtrusiveness. The general idea is that complex patterns will be split into several small patterns with a simple structure, e.g. the sequential sub-activity “fridge usage” is again a weighted indicator for the fuzzy ADL “Preparation of Meal”. The rule format allows the creation of self-contained specialized agents for pattern recognition (more information see [9]). Based on the experiences, the set of agents has been increased, so that other functionalities (in special with focus on “reaction”) are covered, too. The agent set is not limited, so that agents with not covered functionalities can be added easily.

The communication of the agents is handled internally, but the communication with the AAL-System is made in the publish-/subscribe-style like the “reasoner” shown in Figure 1.

5 Conclusion and Future Work

We examined the problem of redundancy in the use of multiple concurrent AAL solutions. We have shown that it is possible to conduct sensor data fusion and aggregation on an AAL system, at which a minimal set of sensors and actuators can be employed, exemplifying this with the application on the universAAL platform.

Currently, in the scope of the BMBF-funded project *OptimAAL*, a prototype operating on the universAAL platform is being developed according to the concept we introduced in Section 3. The prototype will first be tested in a laboratory setting that contains a mimicked home. Following a successful test, a field study taking place in real apartments for the elderly is being prepared for validation of the presented approach. The next step in improving the presented system is to take Quality of Service (QoS) into account to rate setups. Other cost functions such as acquisition costs, running costs, space and volume taken up by the devices are possible candidates. Given a user preference, these considerations could increase the automation in the semi-automatic approach.

While so far we have only taken technical restraints into account, it is evident that non-technical restraints, arising from personal, cultural or medical condi-

tions, can render a particular setup unacceptable to a user. Preliminary feedback has shown for instance, that audio signals might be unsuitable for alerting the elderly because of diminished hearing capabilities. Identifying and incorporating these additional constraints requires expertise in a number of different disciplines.

Acknowledgements. The work presented in this paper was carried out in the OptimAAL-Project (Competence Platform for the development and introduction of AAL-Solutions), funded by the German Federal Ministry of Education and Research (BMBF).

References

1. Elmenreich, W.: Sensor Fusion in Time-Triggered Systems. Ph.D. thesis, Vienna University of Technology, Vienna, Austria (2002)
2. Frikken, K.B., Dougherty, J.A.: An efficient integrity-preserving scheme for hierarchical sensor aggregation. In: Proceedings of the First ACM Conference on Wireless Network Security (WiSec 2008), pp. 68–76. ACM, New York (2008)
3. Katz, S., Ford, A.B., Moskowitz, R.W., Jackson, B.A., Jaffe, M.W.: Studies of illness in the aged: The index of ADL: A standardized measure of biological and psychosocial function. *Journal of the American Medical Association* 185(12), 914–919 (1963)
4. Manulis, M., Schwenk, J.: Security model and framework for information aggregation in sensor networks. *ACM Trans. Sen. Netw.* 5(2), Article 13 (2009)
5. OASIS - Open architecture for Accessible Services Integration and Standardization: OASIS common hyper-ontological framework, COF (2009), http://www.oasis-project.eu/docs/OFFICIAL_DELIVERABLES/SP1/D1.2.1/OASIS-D121.pdf (online - accessed November 2011)
6. Open Geospatial Consortium Inc.: OGC Sensor Web Enablement: Overview and High Level Architecture (2007)
7. Open Geospatial Consortium Inc.: OpenGIS Sensor Model Language (SensorML): Implementation specification (2007)
8. Prückner, S., Madler, C., Beyer, D., Berger, M., Kleinberger, T., Becker, M.: Emergency Monitoring and Prevention - EU Project EMERGE. In: Deutscher AAL Kongress 2008, January 29-31, pp. 167–172. VDE-Verlag, Berlin (2008)
9. Storf, H., Kleinberger, T., Becker, M., Schmitt, M., Bomarius, F., Prueckner, S.: An event-driven approach to activity recognition in ambient assisted living. In: Tschelligi, M., de Ruyter, B., Markopoulos, P., Wichert, R., Mirlacher, T., Meschterjakov, A., Reitberger, W. (eds.) *Aml 2009*. LNCS, vol. 5859, pp. 123–132. Springer, Heidelberg (2009)
10. Universal Open Architecture and Platform for Ambient Assisted Living: The universAAL reference architecture (2010), <http://universaal.org/images/stories/deliverables/D1.3-B.pdf> (online - accessed November 2011)
11. W3C: Resource Description Framework (RDF), <http://www.w3.org/RDF/> (online - accessed November 2011)
12. W3C: OWL-S: Semantic Markup for Web Services. Service Profiles (2004), <http://www.w3.org/Submission/OWL-S/> (online - accessed November 2011)