C-ANIS: a Contextual, Automatic and Dynamic Service-Oriented Integration Framework

Noha Ibrahim, Frédéric Le Mouël and Stéphane Frénot
{noha.ibrahim, frederic.le-mouel, stephane.frenot}@insa-lyon.fr

ARES INRIA / CITI, INSA-Lyon, F-69621, France

Abstract. Ubiquitous computing environments are highly dynamic by nature. Services provided by different devices can appear and disappear as, for example, devices join and leave these environments. This article contributes to the handling of this dynamicity by discussing service integration in the context of service-oriented architectures. We propose C-ANIS: a Contextual, Automatic and dyNamic Integration framework of Services. C-ANIS distinguishes two different approaches to service integration. Automatic integration automatically extends the capabilities of an existing service, leaving the interface unchanged. On-demand integration builds a new service on request from a list of given services. We have implemented C-ANIS based on the OSGi/Felix framework and thereby demonstrated the feasibility of these two service integration concepts. We have also implemented a toolkit providing two different techniques: Redirection, i.e. calling interfaces and replication, i.e. copying implementations of services¹.

1 Introduction

Ubiquitous computing environments are highly dynamic by nature. Services provided by different devices can appear and disappear as, for example, devices join and leave these environments. This article contributes to the handling of this dynamicity by discussing service integration in the context of service-oriented architectures. We propose to distinguish two different approaches to service integration: automatic integration and on-demand integration. Automatic integration automatically extends the functionality of an existing service *S* by integrating it with compatible services in the environment, but leaving the interface of *S* unchanged. This way, the extension in functionality of *S* can be kept transparent to applications or users employing this service. On-demand integration builds a new service on request from a list of given services. It integrates an existing service *S* with a list of services in the environment, creating new interfaces. This integration is initiated by users or applications, and the new interfaces are employed by these latter.

¹ This work is part of the ongoing European project: IST Amigo-Ambient Intelligence for the Networked Home Environment [1].

In line with the service paradigm, we assume that every relevant context parameter of a ubiquitous computing environment is provided by some service. Consequently, we generally define context as the collection of services available in such an environment. Employing this definition of context, we propose C-ANIS: a Contextual Automatic and dyNamic Integration framework of services. C-ANIS integrates automatically and on-demand the available services at run time while taking the whole context into account, and if intended, it can also dis-integrate services again.

A use case is described all along the article to motivate, explain and evaluate our two integration approaches.

Fig. 1. use case

The use case defines three services:

- The webcam service: a service that enables to take a photo via a webcam.
- The storage service: a service that enables to store an object on a device. Two different services offer the same functionality. One implementation is for local storage, the other one for remote storage.
- The naming service: a service that execute a naming strategy defined by a user to name his files and objects.

These services are provided by different devices (cf. fig 1) that can join or leave the environment leading these services to appear and disappear at any time. We have implemented C-ANIS based on the OSGi/Felix framework and thereby demonstrated the feasibility of the two service integration concepts. We have also implemented a toolkit providing two different techniques to realize the automatic and on-demand service integration concepts: Redirection, i.e. calling interfaces and replication, i.e. copying implementations of services.

In the following, we will start by introducing our service model along with our notion of service integration (section 2). This is followed, section 3 and section 4, by the presentation of our two services integration approaches along with their

life cycle. In section 5, we discuss the implementation of our concepts, followed by a first evaluation (section 6). In section 7, we will review relevant related work to position our work. Finally, we present conclusions and open issues (section 8).

2 Integration in Service-Oriented Architecture

2.1 Service Model

A service is composed of three parts:

- interfaces: A service can hold two kinds of interfaces. Provided functional interfaces defining the functional behavior of the service. Required interfaces specifying required functionalities from other services. A functional interface specifies methods that can be performed on the service.
- implementations: Implementations realize the functionality expected from the service. These are the implementations of the methods defined in the functional interfaces.
- properties: a service will register its interfaces under certain properties. The property is used by the framework to choose services that offer the same interfaces, but different implementations.

We model a functional interface of a service S , its implementation and property as follow:

$$Ifc_S \left\{ \begin{array}{l} m1(params1) \rightarrow r1 \\ \vdots \\ mk(paramsk) \rightarrow rk \end{array} \right.$$

$$Impl_S \left\{ \begin{array}{l} Impl1(m1) \\ \vdots \\ Implk(mk) \end{array} \right.$$

$$property_S : (Ifc_S)_{atomic}$$

Where Ifc_S is one functional interface of the service S , mk the method name, $paramsk$ the list of parameters, rk the return result, and $impl_S(mk)$ the implementation of method mk .

Use case. the use case' services (webcam, storage and naming) are modeled as follows:

$$webcam \left\{ \begin{array}{l} Ifc_{webcam} : getSnapShot() \rightarrow Image \\ Impl_{webcam} : impl(getSnapShot) \\ prop_{webcam} : webcam_{atomic} \end{array} \right.$$

$$\begin{aligned}
\text{storage} & \left\{ \begin{array}{l} Ifc_{storage} : save(Object\ obj, String\ ID) \rightarrow void \\ Impl_{storage} : impl_{local}(save) \\ prop_{storage} : storagelocal_{atomic} \end{array} \right. \\
\text{storage} & \left\{ \begin{array}{l} Ifc_{storage} : save(Object\ obj, String\ ID) \rightarrow void \\ Impl_{storage} : impl_{ftp}(save) \\ prop_{storage} : storageftp_{atomic} \end{array} \right. \\
\text{naming} & \left\{ \begin{array}{l} Ifc_{naming} : getNextName(String\ ID) \rightarrow String \\ Impl_{naming} : impl(getNextName) \\ prop_{naming} : naming_{atomic} \end{array} \right.
\end{aligned}$$

The property describes the interface implementation and specifies whether this implementation is atomic or integrated (resulting from integration). To execute a service, the framework can choose services' interfaces considering the property they publish. If no property is specified the framework will randomly choose a service' interface implementation.

Two services are considered by users/applications to be the same if they have the same functional interfaces. They indeed provide, externally, the same functionalities. The two storage services are considered to be the same by users. The implementations of these services is kept transparent from the users/applications (cf. fig 2).

Two services are considered by the run-time framework to be the same, if they have not only the same interface but especially the same property. Two services publishing the same interface but under different properties are considered by the framework to be different. The properties describe the implementation of the functional interface and different implementations mean different services. For the run-time framework, the two storage interface are registered under two different properties ($storageftp_{atomic}$ and $storagelocal_{atomic}$) and considered as two different services (cf. fig 2).

Our service model is independent of any implementations and can be applied to EJBs [2], Fractal components [3], OSGi bundles/services [4] or Web Services [5].

2.2 Services integration approaches

In ubiquitous computing environments, services provided by different devices can appear and disappear as devices join and leave these environments. These services are employed by users or applications being in the environment. New services only come from new devices joining the environment. The only other

Fig. 2. Different services

way to offer new services in these environment is to respond to an external demand of integration. If new services are offered, without being requested, they are likely not to be used.

For that, we distinguish on demand integration that builds new services upon users/applications' requests and automatic integration that extend the functionalities of existing services.

- On-demand service integration: The framework responds to an external demand by providing new services in the environment. This demand comes from users or applications being in the environment. Applications or users tend to use services available everywhere in the context and would like to, whenever it is possible and/or needed, integrate services offered by the context. In particular, if no single service can satisfy the functionality required by the application, combining existing services together should be a possibility in order to fulfill the request [6]. The result of this integration is a new service with new interfaces (new methods), new implementations (new functionalities) and new properties.
- Automatic service integration: The framework selects automatically all the compatible services in the environment and integrate them. The result of these integrations is the same services enriched with new functionalities. The service interface and methods do not change, only its functionality and property change. This way, the extension in functionality can be kept transparent to applications or users employing the services. Once new services are in the environment, the framework automatically compares these services to existant services and if compatible services are found, the automatic integration can take place.

3 On-demand integration of services

On-demand integration builds a new service on request from a list of given services. We will first define our compatibility notion, followed by the life-cycle

of our on-demand integration approach. Finally, we show its application on our use case example.

3.1 Definition of compatibility

Two services are compatible if they have two compatible functional interfaces. Two functional interfaces are defined to be compatible if they have at least two compatible methods. Two methods are compatible if the return result of one method is of the same type of one parameter of the other method (cf. fig 3).

Fig. 3. Combining compatible methods: method1 & method2

Based on the compatibility definition, we define the integration of services as the combination, two by two, of all their compatible functional interfaces, and so of all their compatible methods. The combination of method1 and method2 (cf. fig 3) creates a new method1 with new parameters type corresponding to the parameters of method2 and part of method1' parameters.

3.2 On-demand integration life-cycle

The integration of a given list of services into one service is done by the integration of the services two by two. When integrating two services (cf. fig 4), all their methods are listed and only compatible methods are selected. The framework selects the most appropriate service' implementations to create the new service. This selection is context aware and must depend on the users/applications preferences. For now no strategies are defined and the selection is done statically. Once the implementations chosen, the new service is created, with its new interfaces, implementations and properties. The new service is installed, started, monitored and its interfaces published. If services involved in the integration leave the environment, the service newly created, is dis-integrated and a new service is created. For that a contextual selection of new service' implementations is done. In the meanwhile, all the calls to the service are buffered.

Fig. 4. on-demand integration life cycle

3.3 Use case example

An on-demand integration example is the integration of service webcam and storage (cf. fig 5). The two methods `save` and `getSnapshot` are compatible. Indeed, the return result of `getSnapshot` is of type `Image` which inherits `Object` the type of one parameter of `save`. The two methods can be combined as shown fig 5, and a new method `saveGetSnapshot` can be created. To choose the most appropriate services' implementations, a contextual choice is made upon users/applications preferences. If a user has constraint device, he will probably prefer to store the image on a remote computer and for that the framework will choose the ftp storage implementation. If the user has a pda and would like to store the photo on his device, the local storage is selected by the framework. For now, strategies are hard coded and chosen statically.

4 Automatic integration of services

Automatic integration automatically extends the functionality of an existing service *S* by integrating it with compatible services in the environment, but leaving the interface of *S* unchanged. We will first define the modified compatibility definition for automatic integration, followed by its life-cycle. Finally, we show its application on our use case example.

4.1 Definition of condition-compatibility

The notion of compatibility is the same as defined for on-demand integration but with additional condition. The automatic integration must remain transparent to the users and applications. The new method1 must have the same signature as the initial method1 so that it can be employed by applications and for that

Fig. 5. on-demand service webcam and storage integration

some conditions must be fulfilled. Method2 must have only one parameter and of the same type as its return result (cf. fig 6).

Fig. 6. Keeping the same signature as method1

The condition that needs to be satisfied in order to have an automatic integration of services without generating new functional interfaces in the context is:

condition. One of the two methods to combine must have only one parameter and this parameter must have the same type as the return result of the method. Two methods are *condition-compatible* if they are compatible and one of the method verifies *condition*. We define the automatic integration of two services as the combination, two by two, of all their *condition-compatible* methods.

4.2 Automatic integration life-cycle

Fig. 7. automatic integration life cycle

Automatic service integration is applied upon each appearance of new services in the context. The integration is contextual because it is very dependent on the services in the context, automatic because it is done by the framework upon each appearance of new services.

For the run-time framework a new service is a service with new functional interfaces or new properties.

New services appearing: If these services have new interfaces and so new methods, the framework applies the method matching algorithm. This algorithm returns a list of all *condition-compatible* methods. The automatic integration can take place and new services are created (same interfaces, new properties). If the services already exist, the framework do the matching on the property to determine if the services are new in the context, which means new atomic property or new integrated property. In case of new atomic property, the framework verifies if the methods of these services belong to the list of *condition-compatible* methods and if it is the case, automatically integrates these methods and creates new services (same interfaces, new properties). In case of new integrated property, the framework needs to insure that no integration must be done if it involves the same services already integrated. This condition insures the stop of our automatic integration. Indeed, the framework never re-integrates services

that were previously integrated. All the new services are installed, published and monitored.

Services disappearing: The framework needs to dis-integrate the integrated services. The call to these services will be automatically redirected to other available services offering same interfaces but with different properties. This redirection is kept transparent to the users and applications.

4.3 Use case example

New services: storage, naming and webcam are now available in the context (fig. 1). The framework automatically executes the steps defined in the life-cycle (fig. 7).

These services have all new interfaces. The framework lists all the interfaces available in the context. Once the interfaces known, a list M of all their methods is created.

use case. fig 8

Fig. 8. M: list of all available methods in the context

The framework selects all the methods in this list that has the same parameter and result type. This matching will return a list C of the methods that fulfil the *condition* defined section 3.1.

use case. fig 9

Fig. 9. C: list of methods that has the same parameter and result type

The framework verifies the compatibility of all the methods of M to all the methods of C. The result is a list of all *condition-compatible* methods.

use case. fig 10

Fig. 10. compatible methods: save and getNextName

The integrated services resulting from automatic integration are services having the same interfaces but different implementations and properties.
use case. fig 11

Fig. 11. Same methods signature, different implementations

The new services are now available in the context and registered under these new properties:

`storagelocalintegrated(namingatomic)`, `storageftpintegrated(namingatomic)`.

These new services are reconsidered for a possible re-integration by the framework. As the interfaces are not new, the properties are checked and only non previously integrated services are allowed to integrate (cf. Table 1).

	webcam _{atomic}	naming _{atomic}
<code>storagelocal_{integrated}(naming_{atomic})</code>	no	no
<code>storageftp_{integrated}(naming_{atomic})</code>	no	no

Table 1. Property matching

The run-time framework reconsiders these new services for integration, but the property matching algorithm indicates that all the integration possibilities

have been already done (cf. Table 1). The run-time framework considers two interfaces registered under the same property to be the same.

5 Contextual service integration toolkit

We developed a toolkit to C-ANIS framework under Felix/OSGi. The OSGi specifications define a standardized, component oriented, computing environment for networked services. Adding an OSGi Service Platform to a networked device (embedded as well as servers), adds the capability to manage the life cycle of the software components in the device from anywhere in the network. A unit of deployment called bundle offers the services in the framework. We implement our developing framework on Felix which is open source implementation of OSGi framework specification.

The integration call is done by the framework.

- automatic integration call:

```
integrate(context);
```

Listing 1.1. Integrating services of the context

The framework executes this integration call upon each entrance of a new service in the context. The new service is compared to all other services available in the environment.

- on-demand integration: Integrating the two specified services is done via an integration call:

```
integrate(webcam, storage);
```

Listing 1.2. Integrating services S1 and S2

In OSGi, creating the service is done by creating the unit of deployment, called bundle. An OSGi bundle is comprised of Java classes and other resources which together can provide functions, services and packages to other bundles. A bundle is distributed as a JAR file. To create a bundle we need to tackle several needs:

- unit of deployment: a bundle to deploy the new integrated service.
- integration glue (Table 2): The java code that do the technical integration. We provide two different techniques: the redirection or interface call, done via method call and RMI, and the replication or implementations copy done via method call to the local replicated implementations.
- needed libraries: in case of replication, the implementations of the replicated services are needed and added to the bundle.
- services dependencies: the new service will have to verify the dependencies of the services involved in the integration.

Once the service created, it is installed, started and its interfaces registered in the context (listing 1.3).

	unit of deployment	integration glue	needed libraries	services dependencies
Redirection	Bundle (jar)	Method Call or RMI		S1, S2
Replication	Bundle (jar)	Method Call	S1 bundle, S2 bundle	dependencies S1, S2

Table 2. Integration techniques

```
Properties props = new Properties();
props.put("StorageIfc", "Storage-integrated(Naming-atomic)");
context.registerService(
    StorageIfc.class.getName(), serv, props);
```

Listing 1.3. Example of a service registration

The run-time framework monitors all the integrated services. For each change in the context involving the integrated services, the framework stops the services and dis-integrates them. For automatic integration, all the calls are redirected to services publishing the same interfaces but with different properties. For on-demand integration, the calls are buffered while the service is re-created with new services' implementations.

6 Evaluation

To test our prototype we implemented the above described use case employing a Logitech USB webcam (vfw:Microsoft WDM Image Capture (Win32):0), two Dell Latitude D410 laptops (Intel(R) Pentium(R) M, processor 1.73GHz, 0,99Go RAM) running Microsoft Windows XP Professional (version 2002) and Ubuntu 6.06 LTS.

We measured the time of our matching algorithm, service-integration techniques, execution of the services (cf. fig 12) and bundles' size (cf. fig 13).

The time of our integration techniques is about 1 second for integrating two services. One can choose which technique to apply depending on the context. The redirection technique is more appropriate for constraints devices whereas the replication technique is more recommended for integrating services executing on devices that disconnect very often. The contextual choice of the technique will be the subject of another article.

The integrated service has the same execution time as any other atomic service (cf. fig 12).

For n services in the run-time framework, the complexity of our automatic matching algorithm is $O(n)$ upon each entry of a service in the context and $O(n^2)$ if a matching is done between all the services of the context.

The matching algorithm is relatively quick, but the automatic integration time is not scalable for large context. For run-time frameworks with 100 services, if matching only takes 329 ms, the integration time is much slower. Adding to

that the time it takes to get distant access between remote run-time frameworks, one can quickly see the limits of the automaticity in large context.

Fig. 12. Average of a 100 test runs

Fig. 13. Bundle size for on-demand integration of webcam and storage

7 Related Work

Basically, the process of service integration responds to an external and explicit request (e.g. from users or applications) by providing new services in the environment. This integration is known in related work as “business logic of a client” [7], “on-demand basis composition” [8] or “users tasks descriptions” [9]. On the contrary and to the best of our knowledge, automatic service-integration is considered more as an adaptation of the service itself rather than an integration of services. The idea of a framework that extends and shrinks automatically is not very exploited in the literature.

Integrating services by matching their interfaces has already been done in [10,11]. The matching is especially done on semantic description of the parameters input and output. The matching descriptions is usually language-described at an abstract level. A service takes in charge to find the services corresponding to the semantic descriptions and to execute them at run-time.

[7] and [8] classify composition of web services framework for ubiquitous computing. They emphasize three important characteristics for service composition: dynamicity ([12,13]), automaticity and context-awareness ([13,14]) of the composition. While a certain number of works dealt with the dynamicity and context awareness of service-composition, few were interested in providing a real automaticity.

In all these approaches, composition of services is seen more as a coordination of the services invocations rather than a technique of creating a new functionality

in the context. Our approach proposes a framework that do the integration of services at run time taking into consideration only the services available in the context and creating new services in the context. The automatic integration depends strongly on what is available and provides enriched services that exist as long as the technique employed to integrate them is viable. We proposed to apply the characteristics defined in the above classifications to the interface matching and provide an automatic, dynamic and contextual service-integration framework for ubiquitous computing, that integrates services not only on a on-demand basis.

8 Conclusion and Future Works

In this article, we proposed C-ANIS: a Contextual Automatic and dyNamic Integration framework of Services. C-ANIS framework realizes automatically and at runtime the integration of available services in the environment, generating enriched services and new services as described in our two integration approaches. C-ANIS framework automatically applies different techniques of integration to compatible services, discovered in the environment, generating on the fly the same services enriched with new capabilities or new services. We have implemented C-ANIS based on the OSGi/Felix framework and thereby demonstrated the feasibility of these two service integration concepts.

The contributions of C-ANIS are in its:

- automaticity: Each time a new service is in the framework, a possible integration is done. We distinguished two major integration approaches: automatic integration and on-demand integration.
- context-awareness: Both automatic and on-demand integrations are context-aware. For on-demand integration, The choice of the services' implementations is depending on the context. For automatic integration, the services available in the context define the integration to do.
- dynamicity: Once the integration decided (method matching signature done), the choice of the implementations is done at run-time and can changes with the context changes.

The perspectives of our approach are:

- generality of the matching algorithm: If a return type of method2 matches several parameters' types of method1, only one match is taken into consideration. In our use case example, the return type of `getNextName` matches both parameters type of method `save`. Only one combination is made as the properties `storagelocalintegrated(namingatomic)` and `storageftpintegrated(namingatomic)` can be associated to storage interface only once. Indeed, the properties specify that two services are already integrated and two methods can not be combined more than once. To resolve that issue, we want to describe semantically our services. The matching will be done on semantic description and not on methods' signature to take all the cases into consideration.

- interoperability: The offered toolkit is only for java technology. We plan to use Amigo interoperable services [1] and extend our toolkit to .Net.
- context-awareness: We want to define contextual strategies for the run-time framework for choosing the integration techniques (replication or redirection) and services' implementations depending on the context.

References

1. Georgantas, N., ed.: Detailed Design of the Amigo Middleware Core: Service Specification, Interoperable Middleware Core. Deliverable D3.1b, IST Amigo project (2005)
2. Monson-Haefel, R.: Enterprise JavaBeans. O'Reilly & Associates (2000)
3. Bruneton, E.: Developing with Fractal. The ObjectWeb Consortium, France Telecom (R&D). (2004) version 1.0.3.
4. OSGIalliance: About the OSGI service platform. Technical report, OSGI alliance (2004) revision 3.0.
5. Iverson, W.: Real Web services. O'Reilly (2004)
6. Ponnekanti, S.R., Fox, A.: SWORD: A Developer Toolkit for Web Service Composition. In: 11th World Wide Web Conference. (2002) Honolulu, USA.
7. Dustdar, S., Schreiner, W.: A survey on web services composition. *Int. J. Web and Grid Services* **1**(1) (2005) 1–30
8. Alamri, A., Eid, M., Saddik, A.E.: Classification of the state-of-the-art dynamic web services composition techniques. *Int. J. Web and Grid Services* **2**(2) (2006) 148–166
9. Mokhtar, S.B., Georgantas, N., Issarny, V.: Ad-hoc composition of user tasks in pervasive computing environments. In: 4th international workshop on Software Composition, co-located with ETAPS'05. (2005) Edinburgh, Scotland.
10. Zhou Zhang, J., Jian Yu, S., Kun Ge, X., Wu, G.: Automatic Web Service Composition Based on Service Interface Description. In: International Conference on Internet Computing, CSREA Press (2006) 120–126
11. Kalasapur, S., Kumar, M., Shirazi, B.: Seamless service composition (SeSCo) in pervasive environments. In: International Multimedia Conference, Proceedings of the first ACM international workshop on Multimedia service composition, ACM Press (2005) 11–20
12. Sun, H., Zhou, X., Zou, P.: Research and Implementation of Dynamic Web Services Composition. In: APPT 2003. Volume LNCS 2834., Springer-Verlag Berlin Heidelberg (2003) 457–466
13. Mokhtar, S.B., Georgantas, N., Issarny, V.: COCOA: CONversation-based Service COMposition in PervAsive Computing Environments. In: Proceedings of the IEEE International Conference on Pervasive Services (ICPS'06), IEEE Computer Society (2006)
14. Mrissa, M., Benslimane, D.: Towards a semantic- and context-based approach for composing web services. *Int. J. Web and Grid Services* **1**(3/4) (2005) 268–286