

Supporting Dynamic Service Composition at Runtime based on End-user Requirements

Eduardo Silva, Luís Ferreira Pires, Marten van Sinderen

Centre for Telematics and Information Technology
University of Twente, The Netherlands
P.O. Box 217, 7500 AE Enschede

{e.m.g.silva, l.ferreirapires, m.j.vansinderen}@cs.utwente.nl

Abstract. Network-based software application services are receiving a lot of attention in recent years, as observed in developments as *Internet of Services*, *Software as a Service* and *Cloud Computing*. A service-oriented computing ecosystem is being created where the end-user is having an increasingly more active role in the service creation process. However, supporting end-users in the creation process, at runtime, is a difficult undertaking. Users have different requirements and preferences towards application services, use services in different situations and expect highly abstract mechanisms in the creation process. Furthermore, there are different types of end-users: some can deliver more detailed requirements or can be provided with more advanced request interface, while others can not. To tackle these issues and provide end-users with personalised service delivery, we claim that runtime automated service composition mechanisms are required. In this paper we present the DynamiCoS framework, which aims at supporting the different phases required to provide end-users with automatic service discovery, selection and composition process. In this paper we also present the developed prototype and its evaluation.

1 Introduction

With the Internet becoming ubiquitous, the use of network-based application services is being increasingly adopted and it is expected to grow in the upcoming years [1]. This is being reflected in many technology developments and innovations, such as, for example, the *Internet of Services*, *Cloud Computing* and *Software as a Service* (SaaS). This is leading to the emergence of large sets of services in different domains. At the same time the use of mobile devices with fast data connections is increasing quite rapidly. In [2] it is reported that by 2013 more than 38% of the European population will access the Internet on their mobile device, which is an increase of 300% compared to the current situation.

These developments are allowing and *pushing* new, more adaptive and personalised application services where the end-users play an active role in the process of service creation. Supporting end-users in this kind of service creation process is a complex undertaking. Different users have different preferences and request services in different context situations, which requires different actions to be taken. Furthermore, end-users expect a high-level of abstraction in the service creation process, they are not able to use

very technical and complex tools. Given this, automation has to be provided to support the end-user in the service creation process. We claim that this can be achieved by using semantic-based service composition approaches. Semantic information enables the use of computer agents, which can automatically reason on the services and user specified requirements. This alleviates the end-user from the burden of some of the details and the complexity in the composition process. We define the problem of automatic service composition based on user requirements as *dynamic service composition*. To cope with this problem we propose a framework for dynamic service composition provisioning, called *DynamiCoS*.

The aim of the DynamiCoS framework is to support end-users' automated runtime service composition. To achieve this automated support, DynamiCoS defines ontologies (domain conceptualisations). The framework allows different service developers to publish their semantically annotated services in the framework. These semantic descriptions have to refer to the framework's ontologies. End-users may have different domain or technical knowledge, which implies that their service request interfaces have to be defined accordingly. DynamiCoS tackles this problem by defining a service request that supports different user interfaces. A service request consists of a specification of *goals* the user wants the service to achieve. A *goal* is likewise used to describe services, specifying the activities (or operations) the services can perform. Goals of users and services are specified according to the framework ontologies (representation of the domain of knowledge), this allows matchings to be found, whenever a service realises the user goals.

This paper is further organised as follows: Section 2 characterises different types of end-users based on their knowledge; Section 3 presents DynamiCoS, a framework to support end-users in the composition process; Section 4 provides an overview of related work; and Section 5 provides our conclusions and directions for future work.

2 End-users Knowledge

An end-user denotes a person that makes use of some functionality of a system. In the context of service composition, we consider that an end-user is a person without advanced development skills. End-users may have different characteristics according to their knowledge in the composition process. In this work we group end-users according to their knowledge in the domain of the service being composed and on their technical knowledge on the tooling supporting the service composition process. An end-user may play two roles in our context, creating and/or executing service compositions.

2.1 Domain Knowledge

In the process of service composition the user has to have some knowledge or idea of the service he wants i.e., what does a service do, who provides it, etc. We refer to this knowledge as domain knowledge. Domain knowledge is obtained by learning, advertisement, interaction with the service providers, etc. End-users may have this type of knowledge when they want to use a service in a given application domain, but most of

the times they have limited knowledge and require some interaction with the application domain(s) to acquire further knowledge to be able to make a decision.

For semantic services, the domain knowledge is explicitly defined, via domain conceptualisation, by domain experts in ontologies and consequently used to describe the services in the domain and to find them.

For example, if a user wants to buy a phone online, he usually has an idea of what type of phone and the market, i.e., phone type/brands, price, stores providing it, etc. The application domain may be described as telephone stores, and they describe resources (phones, phone stores, etc) and actions (search phone, buy phone, etc) the user can take. At the end the end-user knowledge of the domain will allow him to decide what decision to take, and possibly compose actions (or services) to realise it.

2.2 Technical Knowledge

Service composition is being used in different domains and applications nowadays, mainly by professional users or developers, which can handle complex tooling and understand the composition process, and details associated with it. For example, many companies nowadays use Web services technology, and apply WSDL (Web Service Description Language) to describe services in the company, and BPEL (Business Process Execution Language) to compose and coordinate services. However, end-users are not expected to know these technical details. A service composition environment for end-users has to provide higher level of abstraction to their users, so that these users without technical background can make use of them.

For example, if the end-user wants to find a phone and then buy it, this may consist of two services (find and buy services). The supporting tooling has to allow the end-user to find suitable services and then help him on the composition process, by automating such process or by suggesting the user possible directions of action. The supporting tooling may depend on the type of domain application. For example, if the application domain is on e-health where a care taker decides a sequence of activities a given patient has to be delivered with (e.g., measure blood pressure and if it is too high send a message to the patient's doctor), the care giver will have knowledge on the domain, i.e., knows the services that are available in the domain. On the other hand, in the example provided above, where an end-user wants to buy a phone, the user may not know the domain, i.e., the composition environment has to deliver the necessary suggestions and support to guide the user towards the creation of the service (composition) he wants.

2.3 Types of End-user

Table 1 defines a possible classification of types of users, based on their domain and technical knowledge.

A *Layman* is a user who does not know in detail the application domain, neither has knowledge on the tooling supporting the composition process. A *Domain Expert* is a user who knows the application domain but does not have knowledge on the technical tooling that supports the service composition process. A *Technical Expert* is a user who has knowledge on a service composition tool, but does not know the details of the application domain. An *Advanced* end-user is a user who has technical knowledge on

Type of End-user	Domain Knowledge	Technical Knowledge
<i>Layman</i>	No	No
<i>Domain Expert</i>	Yes	No
<i>Technical Expert</i>	No	Yes
<i>Advanced</i>	Yes	Yes

Table 1. Types of Users

the tooling supporting the composition process, and furthermore knows the application domain.

By acknowledging these different types of end-users we believe that we can create environments that better suite the requirements of the end-users in a composition process. We expect that alternative end-users types will be identified, being placed in between the types we identifies so far. However, these types already gives us some information that we can use to shape, positions and evaluate composition environments.

At the moment, DynamiCoS addresses mainly the users that have some knowledge about the application domain, i.e., *Domain Experts* and *Advanced* users.

3 DynamiCoS framework

*DynamiCoS*¹ (**D**ynamic **C**omposition of **S**ervices) supports all the phases and stakeholders of the dynamic service composition life-cycle, as discussed in previous work [3]. Fig. 1 shows the DynamiCoS framework architecture. It also indicates (between parenthesis) the technologies used in the implementation of the different framework components in our prototype platform.

Given the complexity of the dynamic service composition life-cycle and its stakeholders' heterogeneity, we have made the core components of the framework technology-independent. In the framework, a service and a service composition are represented as a tuple and a graph, respectively. A service is represented as a seven-tuple $s = \langle ID, I, O, P, E, G, NF \rangle$, where ID is the service identifier, I is the set of service inputs, O is the set of service outputs, P is the set of service preconditions, E is the set of service effects, G is the set of goals the service realises, NF is the set of service non-functional properties and constraint values. We assume in this work that services are of type *request-response*, i.e., they consist of one activity or operation. A service composition is represented as a directed graph $G = (N, E)$. Graph nodes N represent services, i.e., each node $n^i \in N$ represents a discovered service s^i . A node can have multiple ingoing and outgoing edges. Graph edges E represent the coupling between a service output/effect (or a user input for the services) and a service input/precondition, i.e., $e_{i \rightarrow j} = n_{O/E}^i \rightarrow n_{I/P}^j$, where $i \neq j$ (a service cannot be coupled with itself).

3.1 DynamiCoS modules

DynamiCoS consists of the following modules:

¹ <http://dynamicos.sourceforge.net>

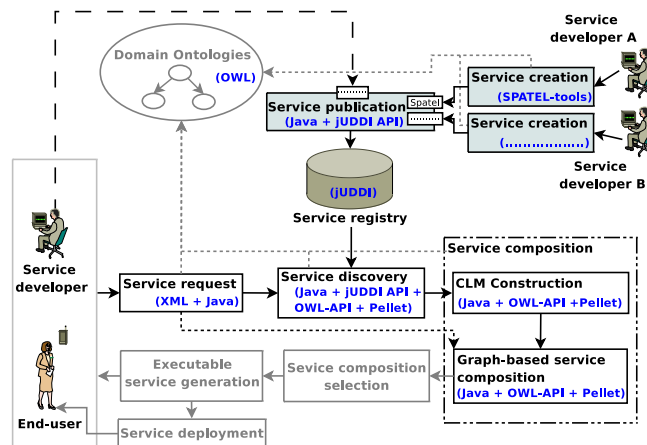


Fig. 1. DynamiCoS framework

Service creation The creation of a new service “from scratch” by a service developer is performed outside the DynamiCoS framework. However, to comply with the capabilities of the DynamiCoS framework, the services have to be semantically described, in terms of inputs, outputs, preconditions, effects (*IOPE*), goals (*G*) and non-functional properties (*NF*), using the framework domain ontologies’ semantic concepts.

Service publication To support the publication of services described in different languages, the DynamiCoS framework has a two-step publication mechanism. First, there should be an interpreter for each supported service description language. The interpreter reads the service description document and extracts the necessary information for publication (*ID, I, O, P, E, G, NF*). This makes the service representation in the framework *language-independent*. Second, the extracted service information is published in the service registry using the DynamiCoS generic service publication mechanism. The service registry allows one to publish, store and discover semantic services.

Service request The user interface to define the service request may have different forms, as long as it gathers the user goals, the expected outputs/effects. Optionally it can also gather the inputs/preconditions the user can provide and non-functional properties he expects for the service. The number of parameters defined on the service request may depend on the type of end-user expertise. If the end-user has technical knowledge and domain knowledge all the parameters (*G, IOPE, NF*) may be used. However, if the end-user has no technical knowledge, a simpler request may be created, for example, only based on the goals the user wants to achieve and expected service outputs. The service request parameters are defined as references to semantic concepts available in the framework ontologies. Therefore, the service request consists of a set of semantic annotations (*I, O, P, E, G, NF*) that describe *declaratively* the desired service properties.

Service discovery The discovery of candidate component services is performed before the service composition phase. Service discovery is based on the service request parameters. The service discovery process consists of querying the service registry for all the services that *semantically* match the defined service request parameters. Since DynamiCoS uses a semantics-based service discovery and composition, it discovers exact matches with the service request G and $IOPE$ semantic concepts, but also other service with partial semantic matches, e.g., services with parameters that are semantically subsumed by the service request parameter concepts, i.e., $RequestedConcept \sqsupseteq DiscoveredConcept$.

Service composition To perform service composition, DynamiCoS first processes the set of discovered services and organises them in a so called *Causal Link Matrix (CLM)* [4]. The CLM stores all possible semantic connections, or *causal links*, between the discovered services input and output concepts. CLM rows (Equation 1) represent the discovered services input concepts ($DiscServs_I$). CLM columns (Equation 2) represent service inputs concepts plus requested service outputs ($ServReq_O$).

$$CLM_{rows} = DiscServs_I \quad (1)$$

$$CLM_{colu} = DiscServs_I + ServReq_O \setminus (DiscServs_I \cap ServReq_O) \quad (2)$$

We place a service s in the row i and column j position if the service has an input semantically related with the input i of the CLM and an output semantically related with the semantic concept on column j . Furthermore, we store the *semantic similarity* of the service output with the column semantic concept and the non-functional properties of the service.

Provided with the CLM, the composition algorithm has to find a composition of services that fulfil the service request. Our service composition algorithm is graph-based. Algorithm 1 presents the algorithm in simplified pseudo code. The composition process starts by analysing the CLM matrix, checking if it contains the requested IOPE. After this, the CLM is inspected for services that provide the service request outputs. If there are services that provide the service request outputs, the algorithm starts by creating the initial matching nodes, otherwise it stops. If the service request outputs can be provided by the discovered services, the algorithm proceeds with a backwards composition strategy towards the service requested inputs, and/or the realisation of all the requested goals. An *open* (or not yet composed) input of the graph is resolved at each algorithm iteration. The algorithm matches the *open* inputs of the services in the graph with the output concepts of services from the CLM matrix, or column concepts. If multiple services exist that match a given service input, a new composition graph is created, representing an alternative service composition behaviour. The algorithm finishes when all requested inputs, preconditions and goals from all the alternative service compositions are resolved.

3.2 Prototype

In our prototype we have used a language called Spatel [5], which has been created in the European IST-SPICE project [6], where the DynamiCoS framework was partly

Algorithm 1: Graph Composition Algorithm

```
Input:  $CLM, ServReq$   
Result:  $ValidComps$   
// Variables  
1  $activeG$ ; // Graph that is active in the algorithm iteration  
2  $activeN$ ; // Node that is active in the algorithm iteration  
3  $openG$ ; // Set of open graphs  
4  $validG$ ; // Set of completed and valid graphs  
// Initialisation  
5 if  $CLM_{rows \cup colu} \supseteq ServReq_{I,O}$  then  
   // Create new graph instantiating the initial Node  
6    $activeG \leftarrow createNewGraph(ServReq)$ ;  
7    $createInitialNodes()$ ;  
8    $openG \leftarrow activeG$ ;  
9 else  
   // Discovered services cannot fulfil the service request  
10  Stop;  
// Graph construction cycle  
11 while  $|openG| > 0$  do  
   // Close graph if it matches  $ServReq_{I,G}$   
12   if  $activeG_{I,G} \supseteq ServReq_{I,G}$  then  
13      $validG \leftarrow activeG$ ;  
14      $openG \leftarrow openG \setminus activeG$ ;  
15      $activeG \leftarrow openG^0$ ;  
16      $activeN \leftarrow activeG_{openN^0}$ ;  
17     break; // Goes to next  $openG$   
   // Checks CLM for services that match open inputs  
18   foreach  $semCon \in activeN_I$  do  
19     if  $CLM_{colu} \supseteq semCon$  then  
20        $activeN \leftarrow CLM_{matchingNode}$ ;  
21     else  
22        $openG \leftarrow openG \setminus activeG$ ;  
23       break; // No possible composition, goes to next open graph  
   // Check if graph NF props comply with requested NF props  
24   if  $activeG_{NF} \cap ServReq_{NF} = \emptyset$  then  
25      $openG \leftarrow openG \setminus activeG$ ;  
26     break; // If Not, composition is not possible  
   // prepare next cycle  
27    $openN \leftarrow openN \setminus activeN$ ;
```

developed. The ontologies used in the framework are described in OWL [7]. We used four ontologies: *Goals.owl*, which contains the services' supported goals and also goals that the user can specify in the service request; *NonFunctional.owl*, which defines framework non-functional properties for services; *Core.owl* and *IOTypes.owl*, which are used to describe services' *IOPE* parameters. However, the framework is general enough to support the use of other ontologies to define other application domains.

For service publication we have implemented a Spatel interpreter. The interpreter imports the Spatel service description by using a Java API generated from the Spatel Ecore model with the Eclipse Modelling Framework (EMF). The service is then published in a UDDI-based service registry that has been extended with semantic support. We use jUDDI [8] as service registry implementation, which is a Java-based implementation of the UDDI specification for Web services. jUDDI offers an API for publication and discovery of services. We have extended the basic jUDDI with a set of UDDI mod-

els (*tModels*) to store the set of semantic annotations (I, O, P, E, G, NF) that describe a service in our framework.

We have created two interfaces to specify the service request: one is a simple Java-based graphical interface, and another is a web-based interface, which allows to experiment with DynamiCoS more easily. Both allow the specification of the different parameters ($IOPE, G, NF$) of the service request. The information introduced by the end-user is then transformed to an XML-based document that represents the *desired service* and sent to the composition framework.

The service discovery is performed based on the service request XML document. The $IOPE$ and G annotations are extracted from this document and the service registry is queried through the jUDDI API Inquiry function. To specify and discover semantically related concepts/services we use the OWL-API [9] and Pellet [10].

The CLM matrix is constructed by using the OWL-API [9], which allows one to handle and perform semantic inference in OWL ontologies by using a semantic reasoner, in our case Pellet [10]. The service composition algorithm is implemented in Java.

In the project website (<http://dynamicsos.sourceforge.net>) we provide more information about the framework and prototype. Furthermore, we provide a web interface for DynamiCoS, where service requests can be issued and the resulting service compositions printed. We provide also some use cases to show the DynamiCoS.

3.3 Evaluation and Validation

In a forthcoming publication we present details on the evaluation and validation of DynamiCoS. The performed evaluations focus mainly on performance and feasibility of the dynamic service composition process. We have concluded that the service composition process can be automated in several phases, by using semantic descriptions. However, semantic reasoning is expensive in terms of processing time. Despite that, we consider that such expensive processing times are realistic to support end-users in the creation of new service compositions on demand at runtime, since in other situations, mainly manual composition is used to tackle this problem, i.e., the user has to spend normally much more time performing a service composition, and we expect manual composition, even with intuitive interfaces, to be too complex for most of the types of end-users.

4 Related work

Dynamic service composition has received a lot of attention lately. We refer to [11] for an overview of some existing approaches. These approaches focus on only some of the dynamic service composition life-cycle phases, mainly the discovery and composition phases. However, some approaches cover most of the phases of the dynamic service composition life-cycle, as, e.g., METEOR-S [12]. METEOR-S provides mechanisms for semantic annotation of existing services, service discovery, and service composition. METEOR-S focuses mainly on design-time creation of service compositions, by developing *templates* that can have dynamic bindings at runtime. Our approach, as many

other ones, has been inspired by METEOR-S, but we target an *on demand* runtime service composition creation, to support end-users at runtime. Kona et al. [13] propose an approach oriented to automatic composition of semantic web services. Similarly to DynamiCoS, they propose a graph-based service composition algorithm. The composition process is performed using a *multi-step narrowing algorithm*. The user specifies a service request, or a *query service*, specifying the *IOPE* for the desired service. The composition problem is then addressed as a discovery problem.

Recently new approaches have emerged to support specifically end-users in the service composition process. These approaches, for example [14] [15] [16] [17], mainly focus on using techniques for mashup, with intuitive graphical representations that allow end-users to create their own services. We argue that these approaches are applicable for some type of users. If the user of the composition environment has some technical knowledge on the composition environment, has a clear idea of the service he wants and knows the application domain, these environments may be appropriate. However, if the end-user does not have a clear idea of the service he wants, but only some vague ideas about goals that he wants to be fulfilled by the service, an approach similar to DynamiCoS may be more appropriate.

5 Final Remarks

In this paper we start with an overview of possible existing types of end-users. End-users may have different knowledge on the application domain, and also on the technical tooling supporting the service composition process. Given this, we claim that different types of supporting environments have to be created or adapted, according the type of end-user and his knowledge/expertise. We propose DynamiCoS, which is an approach that is mainly targeted to users with application domain knowledge, i.e., *Domain Experts* and *Advanced users*. To achieve automation in the composition process, DynamiCoS is based on semantic services. DynamiCoS is neutral with respect to the service description languages used by the service developers. DynamiCoS supports service creation and publication by service developers at design-time, which populate the universe of available services; and automatic service composition by end-users at runtime. We have experimented with DynamiCoS, showing the it is capable of providing *real-time* service delivery, and we observed that semantic reasoning is still an expensive task in terms of processing time. However, these expenses may be worth to pay for, in case automated support of end-users in the creation of their own services is enabled.

In the future we will investigate how the user context and preferences can be used on the optimisation and personalisation to the composition process. Furthermore, we will improve the user interface for service composition, namely by offering to end-users require mechanisms to *guide* them though the process of specification of the service composition behaviour. This will enable the support to all the types of users presented in Section 2.3, specially the ones without domain knowledge neither technical knowledge. This process will require several interactions between the platform and the user and a dynamic negotiation to match the user interests with a composition created out of the available services. The generated service compositions need to be translated to an executable formalism, enabling a full runtime environment to be built. We will address

the translation of service composition representation to an executable language, such as, e.g., WS-BPEL. We plan to support other service description languages, such as, for example, SAWSDL and OWL-S, to create more elaborate service composition scenarios. This will allow us to perform empirical studies on dynamic service composition based on end-users' requirements.

References

1. Gartner: Gartner highlights key predictions for it organisations and users in 2008 and beyond (January 2008)
2. Forrester: European mobile forecast: 2008 to 2013 (March 2008)
3. Silva, E., López, J.M., Ferreira Pires, L., van Sinderen, M.J.: Defining and prototyping a life-cycle for dynamic service composition. In: International Workshop on Architectures, Concepts and Technologies for Service Oriented Computing, Portugal (July 2008) 79–90
4. Lécué, F., Léger, A.: A formal model for semantic web service composition. In: International Semantic Web Conference. LNCS, vol. 4273 (2006) 385–398
5. Almeida, J.P., Baravaglio, A., Belaunde, M., Falcarin, P., Kovacs, E.: Service creation in the SPICE service platform. In: Wireless World Research Forum meeting on "Serving and Managing users in a heterogeneous environment". (November 2006)
6. Cordier, C., Carrez, F., van Kranenburg, H., Licciardi, C., van der Meer, J., Spedalieri, A., Rouzic, J.P.L.: Addressing the challenges of beyond 3G service delivery: the SPICE platform. In: International Workshop on Applications and Services in Wireless Networks. (2006)
7. Smith, M.K., McGuinness, D., Volz, R., Welty, C.: Web Ontology Language (OWL) guide, version 1.0. W3C. (2002)
8. Apache: Apache juddi. <http://ws.apache.org/juddi/>
9. Bechhofer, S., Volz, R., Lord, P.: Cooking the semantic web with the owl api. In: International Semantic Web Conference. (October 2003) 659–675
10. Sirin, E., Parsia, B., Grau, B.C., Kalyanpur, A., Katz, Y.: Pellet: A practical owl-dl reasoner. *Web Semantics: Science, Services and Agents on the World Wide Web* 5(2) (June 2007) 51–53
11. Rao, J., Su, X.: A survey of automated web service composition methods. In: International Workshop on Semantic Web Services and Web Process Composition. (2005) 43–54
12. Verma, K., Gomadam, K., Sheth, A.P., Miller, J.A., Wu, Z.: The meteor-s approach for configuring and executing dynamic web processes. Technical report, University of Georgia, Athens (June 2005)
13. Kona, S., Bansal, A., Gupta, G.: Automatic composition of semanticweb services. In: International Conference on Web Services. (2007) 150–158
14. Liu, X., Huang, G., Mei, H.: A user-oriented approach to automated service composition. In: *Web Services, 2008. ICWS '08. IEEE International Conference on.* (Sept. 2008) 773–776
15. Yelmo, J.C., Trapero, R., Álamo, J.M., Sienel, J., Drewniok, M., Ordás, I., Mccallum, K.: User-driven service lifecycle management — adopting internet paradigms in telecom services. In: *ICSOC '07: Proceedings of the 5th international conference on Service-Oriented Computing*, Berlin, Heidelberg, Springer-Verlag (2007) 342–352
16. Ro, A., Xia, L.S.Y., Paik, H.Y., Chon, C.H.: Bill organiser portal: A case study on end-user composition. In: *WISE '08: Proceedings of the 2008 international workshops on Web Information Systems Engineering*, Berlin, Heidelberg, Springer-Verlag (2008) 152–161
17. Nestler, T.: Towards a mashup-driven end-user programming of soa-based applications. In: *iiWAS '08: Proceedings of the 10th International Conference on Information Integration and Web-based Applications & Services*, New York, NY, USA, ACM (2008) 551–554