

Challenges and Techniques on the Road to Dynamically Compose Web Services

Matthias Fluegge
Fraunhofer FOKUS
Institute for Open Communication Systems
Kaiserin-Augusta-Allee 31 - 10589
Berlin - Germany
fluegge@fokus.fraunhofer.de

Neil Paiva Tizzo
Computer Science Dept.
PUC-Minas University
Av. Fracis Cletus Cox, 1661 - 37701-355
Poços de Caldas, MG - Brazil
neil@pucpcaldas.br

Ivo J. Garcia dos Santos
Institute of Computing
University of Campinas - UNICAMP
P.O. Box 6176 - 13084-971
Campinas, SP - Brazil
ivo@ic.unicamp.br

Edmundo R. M. Madeira
Institute of Computing
University of Campinas UNICAMP
P.O. Box 6176 - 13084-971
Campinas, SP - Brazil
edmundo@ic.unicamp.br

ABSTRACT

One of the great challenges to be faced in order to enable the success of future Web-based applications is to find effective ways to handle with the interoperability demands. In this context, *Service-Oriented Architectures* and *Web Services* technology are being considered as the most affordable solution to promote *interoperability*, by applying strategies like *Service Composition*. Nevertheless, most composition approaches applied nowadays in real world contexts lack dynamism. In fact, there is not yet a consensus regarding what would really be a *dynamic composition*. In this paper we propose some criteria to identify the levels of dynamism and automatization in service compositions. Furthermore, taking into account a model driven approach, we propose a strategy where different techniques can be used to make compositions more dynamic and automatic. This strategy is then exemplified and discussed considering an e-Government composition scenario.

Categories and Subject Descriptors

H.3.5 [Information Storage and Retrieval]: Online Information Services—*Web-based services*; D.2.12 [Software Engineering]: Interoperability; I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search—*Plan execution, formation, and generation*

General Terms

Design

Keywords

Web Services Composition, Semantics, Service Selection, Automatic Plan Generation, Model Driven Architecture

1. INTRODUCTION

The term *interoperability* can be defined as 'the ability of two or more systems or components to exchange information and to use the information that has been exchanged' [11]. In fact, interoperability plays a key role in the new world of networked applications, specially in the e-Business and e-Government domains. In this context, much has been discussed in the literature regarding SOA (*Service Oriented Architectures*) and the so-called *dynamic* (or *automatic*, *adaptive* and even *autonomic*) *Service Composition*, but the fact is that there is up to the moment no consensus regarding the exact definition and broadness of these terms. The first contribution of this paper is the proposal of a set of parameters to classify the compositions into different levels of dynamism and automation. Based on these parameters, we then identify that most composition approaches applied in real-world contexts using traditional Web service technologies [14] (such as WSDL and BPEL) can be classified as *static* since the process model is created manually and the services are bound at design time. There are some attempts to apply late binding of services based on fixed interfaces and message formats, being these considered to be semi-dynamic service compositions. Approaches exposing a higher degree of dynamics can hardly be found in a real world context. Apart from obstacles such as performance and trust, the reason is clearly related to the fact that *traditional* Web services just partially kept their promise of being self-contained and self-describing software components.

Taking all this into account, we propose a path to increase the levels of dynamism and automatization in the service composition process, showing where resources such as ontologies and Artificial Intelligence (AI) techniques could be applied using the Model Driven Architecture (MDA) approach. The strategy is then exemplified through a composition in the e-Government context. Throughout the paper, our main goal is to discuss more general strategies and techniques, which can be further applied in different scenarios, instead of focusing on some specific technology or implementation. This article is organized as follows: Section 2 presents an overview of the steps of a service composition process; Section 3 compares different composition strategies and draws one path towards a fully dynamic composition process; Section 4 presents an example in the e-Government context; and finally in Section 5 conclusions and final remarks are stated.

2. THE SERVICE COMPOSITION

A composite service can be regarded as a combination of activities (which may be either atomic or composite services), invoked in a predefined order and executed as a whole. In this way, a complex service has the behavior of a typical business process. In order to build a service composition, some steps must be taken (not necessarily in this order): (1) A *process model* specifying control and data flow among the activities has to be created; (2) Concrete services to be bound to the process activities need to be discovered. The service composer usually interacts with a broker, e.g. a service registry, in order to look up services which match with certain criteria; (3) The composite service must be made available to potential clients. Again the broker is used to publish a description and the physical access point of the service; (4) During invocation of a composite service a coordinating entity (e.g. a process execution engine) may manage the control flow and the data flow according to the specified process model (see Section 2.3). Next we further analyze some characteristics of these steps which we will use later as criteria to classify the compositions.

2.1 Discovery, Selection and Binding

The selection of the activities which will participate in a service composition may be done either at design time or at run-time. In the former, the bindings are static, i.e. each instantiation of the composite service will be made up of the same constituent services. In the latter, the constituent services are selected at run-time, based on automatically analyzable criteria, such as service functionality, signature and QoS parameters. Late binding implies the dynamic invocation of the constituent services, i.e. a sufficient level of interoperability has to be established, either through fixed interfaces or by applying more sophisticated matchmaking and mapping mechanisms. For a service provider, the applied binding mechanism has several business implications. In a growing service market, third party service providers may offer the same functionality at different conditions, e.g. regarding QoS parameters like price. Applying late binding, the discovery and invocation may become scalable as the number of services increases. Thus the costs of a composite service offered by a provider may decrease along with the growing competition in the associated marketplace. The cost advantage can be either handed over to the consumer or it will increase profitability at the provider's side. Further-

more, late binding may enhance fault-tolerance and thus reliability. Since the actions in a process are not hardwired to concrete services, the unavailability of a service may be compensated through the invocation of a functionally equivalent one. In addition, there are some scenarios where important service characteristics (like price) change constantly, what makes the use of run-time service discovery almost essential for the success of the composition. On the other hand, in some specific application domains, the lack of determinism, i.e. the fact that it is not possible to previously know which service is going to be selected, is not acceptable.

2.2 Creation of the Process Model

Another significant characteristic of a service composition strategy is the degree of automation in the creation of the process model. Traditional service composition methods require the user to define the data flow and the control flow of a composite service manually, either directly or by means of designer tools, e.g. in a drag-and-drop fashion. Subsequently the process description is deployed in a process execution engine. Depending on the abstraction level provided by the tools and also on the applied binding mechanism, the user either creates the process model based on concrete service descriptions or based on abstract service templates which are representatives for sets of services, i.e. for service classes. With respect to the multitude of available services and service templates, it may be a time-consuming task to manually select reasonable building blocks for the composite service. Furthermore, the creation of the data flow, i.e. the parameter assignments between the activities, can be complex and might require the user to have extensive knowledge about the underlying type representations. More advanced composition strategies actively support the user with the creation of the process model, which is often referred to as *semi-automated service composition*. Corresponding modeling tools may interact with a broker in order to automatically look up services which match (regarding IOPEs - Inputs, Outputs, Preconditions, Effects) with the already available control and data flow, thus facilitating and accelerating the creation of the process model. The same applies for the creation of models that are based on abstract functional building blocks (which will be bound to concrete services at run-time). Parameter assignments between these building blocks may be automatically recommended based on an analysis of the underlying types and concepts.

Fully-automated composition approaches intend to generate a service composition plan without human interaction. Mostly AI inspired methods based on formal logic are used for that matter, such as automated reasoning through theorem proving. By means of a planning algorithm a workflow graph containing available activities or concrete services is generated to satisfy the requirements established by the requestor. If there are multiple solutions for the problem, i.e. several plans satisfy the given set of requirements, a selection is made based on QoS parameters. This selection can either be made by the process designer or automatically through predefined weighting and ranking functions. Combining the latter with late service binding implies that the complete service composition (i.e. process model generation and service selection) can be performed at run-time. The question to which extent the composition procedure can be automated is subject to research. Fully automated service composition may work in narrow and formally well defined

application domains. The more complex the context however the more difficult it will be to apply the automated service composition approach in real-world applications. Again the applied degree of automation for generating the process model has significant business implications for a provider who composes services and delivers them to consumers. As mentioned above, modeling the control flow and the data flow of a composite service may be time-consuming tasks. (Semi-) automated composition techniques promise to speed up this procedure, thus bringing down the costs for developing new services. Furthermore time-to-market is accelerated since the provider may react faster and more flexible to the customer requirements. In addition the designed composite services improve in quality as the application of "intelligent" tools helps to create more efficient processes, e.g. by proposing parallel execution of functionally independent activities. One interesting (and feasible) approach for semi-automated compositions was proposed in the SATINE project [5]. It is based on *self-contained activity components* [9], which are created semi-automatically based on OWL-S [4] service ontologies.

2.3 Execution

When composing Web Services, two different execution models are usually applied: *Orchestration* and *Choreography*. There is not a common sense regarding these two definitions, but we can consider that in an *Orchestration* all interactions that are part of a business process (including the sequence of activities, conditional events, among others) must be described, like on a traditional workflow system. This description is then executed by an orchestration engine, which has control of the overall composition. On the other hand, a *Choreography* is more collaborative and less centralized in nature. Only the public message exchanges are considered relevant and more, each service only knows about its own interactions and behavior. Differently from Orchestration, there is not an entity that has a global view/control of the composition [17, 20]. If we refer to the origin of the words, a good comparison can be made. The first, orchestration, can be compared to a set of musicians (*services*) commanded by a conductor (*engine*). The second, choreography, can be compared to a group of dancers (*services*) that already know how to perform and that don't obey to a central coordination. Usually real scenarios involving complex systems with multi-part interactions demand both approaches. In [6] the authors propose a set of policies to regulate service compositions and establish a relationship among these policies and the execution models.

3. TOWARDS A DYNAMIC SERVICE COMPOSITION PROCESS

In this section we first present a clear classification of different composition strategies in terms of dynamism and automation. Then we draw a possible path towards a fully dynamic composition process based on a model-driven approach.

3.1 Service Composition Strategies

Besides the execution model (orchestration or choreography), two service composition characteristics have been examined, namely the type of service discovery/selection/binding and the degree of automation applied for the creation of

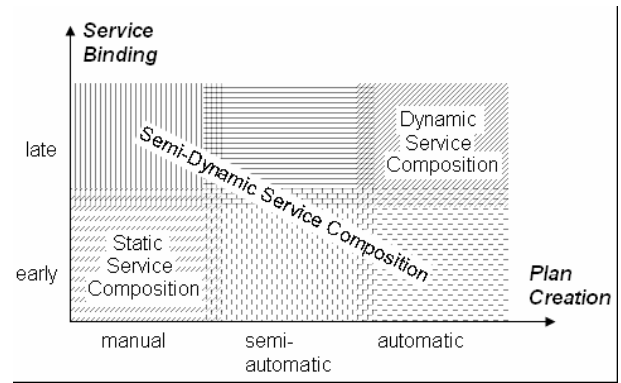


Figure 1: Classification of Service Composition Strategies

a process model: service composition approaches may use early binding or late binding; the process model can be created manually, semi-automatically or automatically. As illustrated in Figure 1, these characteristic values can be used for a classification of existing service composition strategies in six main categories. The fact that the borders between these categories are not strict but fluent is made clear through the smooth transitions between the squares. Some categories may overlap, i.e. there are composition approaches that may be assigned to two or more categories. To give an example: besides early and late binding there may be several variations in between, such as the specification of a restricted set of service candidates at design time from which one service is chosen and invoked at run-time.

There are numerous examples for each of the identified categories. In EFlow [2] and similar systems the data flow and the control flow of a composite service need to be defined manually. The services however can be bound early at design time or at run-time by means of appropriate proxies. An example for semi-automated creation of the process model based on semantic service descriptions has been presented by *Sirin et. al.* [21]. All possible services that match with the current activity are presented to the user. The semi-automated composition approach described by *Fluegge and Tourtchaninova* [9] uses abstract functional building blocks which are bound to concrete services at run-time. The user is supported with the creation of the process model by automatically analyzing the input and output concepts of the building blocks. An example for fully automated service composition based on AI-inspired planning algorithms is given by *Ponmekanti and Fox* [18]. A strategy for automatic composition based on a goal description language (GDL4WSAC) is proposed by *Lin et. al.* [13].

In the previous discussion regarding the implications for an actor who creates and provides composite services, it was argued that composition approaches applying late binding mechanisms are more adaptable to a changing environment, where third party providers are frequently leaving and joining. Furthermore it was argued that a high degree of automation during creation of a process model cuts down development costs and accelerates time-to-market, thus resulting in a higher flexibility of a composite service provider. In addition, quality aspects, such as reliability, have been considered. When combining the terms adaptiveness and flexibility to the more generic term dynamics, a coarse-grained

and more business-oriented classification in static, semi-dynamic and dynamic service composition strategies can be made (see Figure 1). Taking into account the above mentioned attributes cost efficiency, time-to-market and reliability, it can be argued that a high degree of dynamics for service composition has positive effects on the providers' profitability.

On the other hand this does not inherently mean the more automation the better. The degree of dynamics applicable in a real world context is limited by many more factors, being trust (see Section 3.2) one of the most relevant. In addition, performance issues can also represent a problem, since interacting with a broker for service discovery and match-making as well as applying sophisticated AI algorithms for automated plan generation may be time-consuming tasks.

3.2 Trust

Dini [7] states that "functional optimization addresses very well the technological part of a problem, but does not pay enough attention to the intentions of the users and the structure and dynamics of social groups and user communities". Applying dynamic service composition usually results in limited control over the way *how* a problem will be solved, respectively how a service will be composed. As a consequence it might become opaque with which partners the provider of a composite service and its customers collaborate and share data in the scope of a process. Trust is the keyword in this respect and the on-demand establishment of trustable relationships is still an unsolved problem.

3.3 Model Driven Approach

The Model Driven Architecture (MDA)[15] is a new approach proposed by the Object Management Group (OMG) to develop applications and write software specifications. It is based on three standards: the Unified Modeling Language (UML), the MetaObject Facility (MOF) and the Common Warehouse Metamodel (CWM). These standards should facilitate the design, description, storage and exchange of models. The MDA approach separates the specification of the operation of a system from the details of the way that system uses the capabilities of its platform. It uses abstract models to specify all the logic of the application where concepts on languages or platform are irrelevant. Later, these models are used to create new models which express the requirements of the system in a specific platform. Note that *Platform independent* and *platform specific* are not absolute concepts: what is specific to one system can be independent to another. The MDA specifies that the following models should be created during a development process [15]:

- *Computation Independent Model* (CIM): also called a domain model, focuses on the environment and requirements of the system. The details of the structure and processing are hidden;
- *Platform Independent Model* (PIM): provides a description of the system from a platform independent viewpoint, focusing only on the system functionalities;
- *Platform Specific Model* (PSM): describes the system combining the specifications in the PIM with the details regarding the platform where the system will run.

As the use of MDA implies an increase in the abstraction level during system development, the developer, start-

ing from a more abstract level, can successively produce lower abstract models until in the end an executable model is reached. All models must be written in a formal model language, in order to enable the transformation between models to be computer aided. But note that it is not mandatory to use the same language in all models, i.e., it is also possible that a language transformation between different models of the same system becomes necessary. The great advantage in using MDA is the ability to transform a platform independent model (PIM) into a model capable of running into a great variety of technologies. MDA assumes that technology is very volatile, thus an automatic PIM-PSM transformation can save time and money by improving the efficiency of steps like implementation, integration, maintenance, testing and simulation. In an ideal world, the developer would simply submit the PIM to generators which would produce in the end executable code. But the reality is different and we are far away from this - in practice a lot of manual work must still be done.

In a recent work, ISO/IEC [12] defined the use of the UML for expressing system specifications in terms of the five viewpoints defined by the *Reference Model for Open Distributed Processing* (RM-ODP): enterprise, information, computational, engineering and technology. A relationship with RM-ODP viewpoints and MDA can be made as following: the CIM is provided by an enterprise viewpoint with relevant parts of an information viewpoint; the PIM is composed of the information and computational viewpoints; and finally, the PSM is related to the engineering viewpoint. The technology specification does not have an equivalent in the MDA. For each viewpoint there is an associated viewpoint language which can be used to express a specification of the system from that viewpoint, being the model/language transformations provided by the MOF.

In Figure 2 we see a composition process following the MDA approach. The process starts at the definition of CIM, usually a manual task (step 1). This definition must include, among other things, the identification, specification and modeling of the composition. The UMM (UN/CEFACT Modeling Methodology) [3] is an example of methodology that could be used at this phase of the project. The first transformation takes place into CIM-PIM (step 2). The transformation between models can be complex and, in almost every case, parameters need to be set in the source model in order to drive the transformation. After the transformation, refinements should be performed in the PIM in order to go in the direction of the executable code (step 3). In order to stay aligned with the RM-ODP, the information and computational viewpoint languages must be used. Otherwise, BPMN (Business Process Modeling Notation), UML activity diagram or EDOC profiles are also typical languages used to design the PIM in a service composition context.

Next, a transformation from a PIM into a PSM takes place (step 4). This transformation plays a special role in MDA and a mapping language can be defined to automate it. In an ideal world, the same PIM could be transformed in several PSMs, for instance, a PSM-J2EE, PSM-CORBA, PSM-.NET, PSM-WS or any other. The PSM can be specified using the engineering viewpoint language defined by ISO/IEC or UML activity diagrams with extensions for a specific platform. As the previous transformation, some parameters can be set in the PIM to drive the transformation and, after it, the refinement of the PSM should be neces-

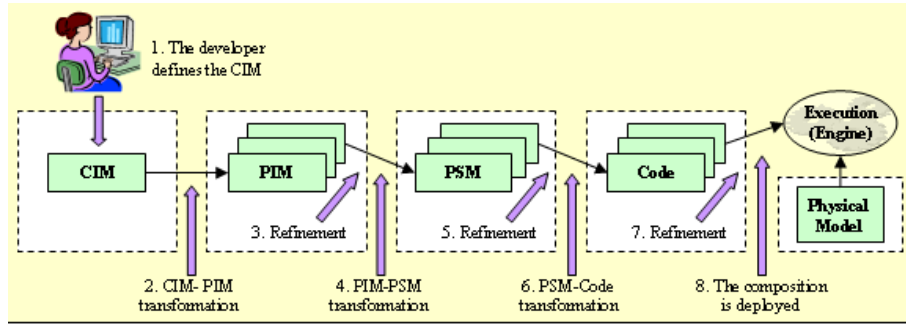


Figure 2: The Composition process following an MDA approach

sary (step 5). *Bzivin et al.* [1] describes a PIM (UML and EDOC) transformation to PSM (Java and Web Services), focusing on the static aspects of the mappings. *Patrascotu* [16] presents the mapping from EDOC profiles to Web Services using a transformation language called YATL, also considering the dynamic aspects of a composition.

A framework proposal for service composition using the MDA approach applied to the e-Government area was presented by *Tizzo et al.* [22].

Finally, the last transformation: PSM-code (step 6). The code is the composition described in some executable language (BPEL for example). A very detailed PSM can describe the whole service composition logic. So, the gap between PSM and code can be very small. But, as the previous transformations, fine adjustments can be necessary before actually running it (step 7). The code, added to the deployment description, completes the models that are necessary to run the composition (step 8). Analyzing the MDA approach in a reverse way, the automatic code generation would start in the PSM-code transformation. When it's possible to produce a full executable code from this transformation, the code would be hidden and in fact the PSM would be executed by a virtual machine. This process is analogous to the one that happens with a traditional compiler or interpreter. Going further, if we apply this automatization within the PIM-PSM transformation, a virtual machine could execute the PIM. The next section presents the challenges and guidelines in order to archive this goal.

3.4 Using Semantics and MDA to Increase the Dynamism in Service Composition

According to the MDA (see previous section), the automation may take place in two different directions: from one model to another (model transformations) and inside a model (model refinement). During these steps, the abstraction level is gradually reduced (see Figure 3). Note that the service compositions, which are described using abstract services, need to be binded to concrete services at a certain moment in time. The use of semantic descriptions and ontologies plays a special role in this stage. The composition described in a CIM is a description of a sequence of tasks and its data and control flow. At this point, a task is an abstract service, i.e., it is only a service description and it may not have an implementation. The information used to describe a task is composed of four fundamental elements: signature, preconditions, post conditions and information invariants [10]. Non-functional aspects can also be described. The activities in a CIM must then be detailed,

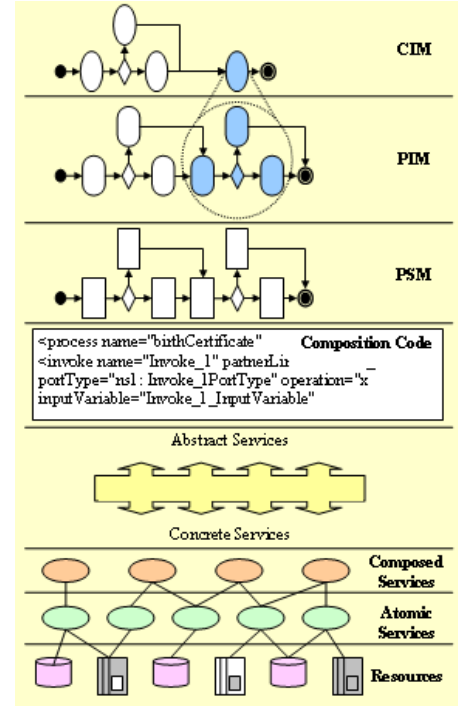


Figure 3: Viewpoints and Abstraction Levels

process done during the CIM-PIM transformation and/or PIM refinement. The automatic (or semi-automatic) transformation and refinement must rely on semantic descriptions to provide machine-readable information about each task. Algorithms based on AI techniques [19] (such as Situation calculus, PDDL, Rule-based planning or Linear Logic) can use these descriptions to decompose CIM tasks or refine the PIM. As already mentioned in Section 2.1, one of the characteristics that can be found in dynamic compositions is the possibility of selecting the participant services. Theoretically, this can be done at the PIM, PSM, code or run-time (late-binding): when it will happen can be determined by a technology restriction or by a project decision: (1) the BPEL language does not allow service selection at run-time, forcing the developer to do it before; (2) on the other hand, when the process environment changes over time, the developer can decide to do the selection as later as possible.

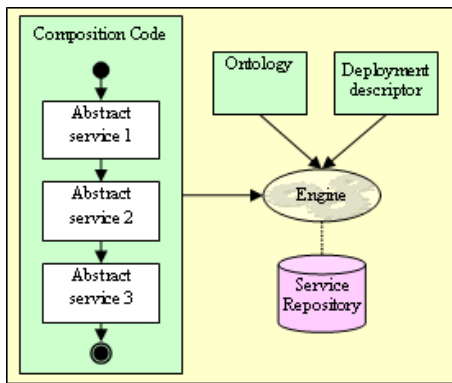


Figure 4: Composition with late (semantic-based) binding

In order to perform the selection of services and considering the late-binding scenario, the result of step 8 (Figure 2) would be a composition of Abstract Services, to be bound to Concrete Services at run-time. In Figure 4, this process is illustrated. An engine receives the composition description (i.e. the code), starts to run it and for each abstract service (described with semantic information), a service repository (or a service broker) is contacted in order to discover which service will be the responsible for executing that activity, always considering the associated ontology.

4. AN EXAMPLE

In order to apply the techniques presented in Section 3, we consider now an example in the e-Government context. The goal of the composition is to provide *birth certificates* to citizens through an e-Government portal. First of all, the portal checks the municipality where the citizen was born (*city of birth*). This is an essential step in order to correctly determine which services will participate in the composition. For different municipalities, even though the CIM is the same, different services should be selected and a different sequence of activities might be performed. That is also an important fact to justify the use of a dynamic strategy in this scenario. Furthermore, each one of the activities involved in the process can be performed by services located in different organizations. This would also create the necessity of considering aspects such as privacy, autonomy and security, omitted in the example for not being in the scope of this paper. An interesting strategy for handling these cross-border issues is the use of *Interaction Policies* [6].

Referring back to the strategies presented in Section 3, and considering a unique CIM as input, the information "city of birth" may be used in one of the following manners:

1. Considering that each activity of the CIM could be implemented in a different way for different cities, the "city of birth" is used to determine the result of the CIM-PIM transformation and refinement. That is, each city might have a different associated PIM;
2. If besides the same CIM, the PIM is also identical for all municipalities, the differences might occur in the PSM. It means that the activities are exactly the same for all cities, but each city may rely on a different technological platform;

3. When the CIM, PIM and PSM are identical, only one composition is built, deployed, and the "city of birth" is used as parameter for selecting the correct concrete services to perform each of the activities in a given execution instance.

Note how one single variable, depending on the strategy adopted and on the characteristics of the problem, may have a great impact on the result composite service.

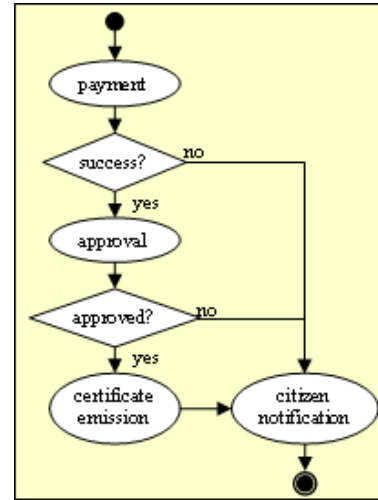


Figure 5: CIM (UML Activity Diagram)

Remember that in our example, the citizen requests the emission of a birth certificate at an e-Government portal. The steps of this process described in the CIM (Figure 5) are:

1. The citizen has to pay for the emission of the certificate;
2. If the payment succeeds, the process continues and the request is forwarded to a public servant (step 3). If not, the citizen must be notified about the problem;
3. A municipality servant must check the pending certificate emission request and validate them. In case of his approval, the request proceeds to step 4. If the servant, for any reason, denies the request, the citizen must be notified;
4. Finally, the certificate must be emitted and the citizen notified about the status of his request.

Other requirements of the process could also be modeled in the CIM (in fact they are necessary if we desire to automatically proceed to next step in the MDA process), but it is not in the scope of this article to further detail them.

As shown in Figure 2, the next steps in our strategy are then the transformation of the CIM into a PIM and the PIM refinement (Figure 6). We consider in our example that each municipality may have its own PIM, and that the information "city of birth" will be important in all transformations and refinements and also in the execution phase (to correctly select and bind the abstract services to concrete implementations). Therefore, the PIM-PSM and PSM-Code transformation would be analogous to the CIM-PIM transformation illustrated in Figure 6.




Figure 6: CIM-PIM transformation and refinement

In order to implement these transformations we consider the associated ontology, the services registry and also a database with previous existing activity models. This model database is particularly important because it contains, at each level, a description of the possible activity mappings to the immediate lower abstraction level in the composition (see Figure 3). Note that if we do not rely on this strategy, we must then adopt some AI technique (Section 2.2) to perform the abstraction level change.

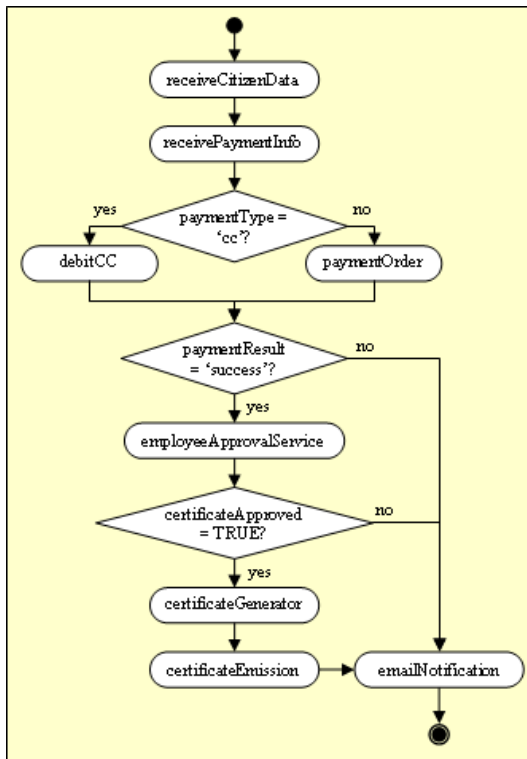


Figure 7: PIM (UML Activity Diagram)

In Figure 7 we see the UML Activity Diagram which is part of the PIM for a given municipality. Differently from the CIM, the PIM describes in more details the sequence of activities that should be performed to successfully complete the process:

1. An activity called "receiveCitizenData" is responsible for getting the citizen request and information;
2. The payment information is sent to the activity "receivePaymentInfo", that checks whether the citizen will pay by Credit Card or by Payment Order, and

then calls the associated activity to actually perform the payment ("debitCC" or "paymentOrder");

3. If the payment was successful, the process proceeds to step 4. If not, an e-mail notification activity must inform the citizen the reason of the failure;
4. A public servant participates in the activity "employeeApprovalService" and checks if everything is fine with the request, validating it. In case of approval, the process proceeds to step 5. In case of failure, the citizen must be informed by the activity "emailNotification";
5. The certificate is generated electronically ("certificateGenerator") and then printed and stamped ("certificateEmission"), being finally the citizen notified of the success of his request and informed about the procedures to pick up the document.

Next follows a (possible automatic) transformation from the PIM into a PSM which can be, for instance, a BPEL specific one. Finally a PSM-Code transformation takes place, and the resulting composition code (made of abstract services) is deployed. At run-time, the binding to the concrete services is performed like Figure 4 illustrated.

5. CONCLUSIONS AND FINAL REMARKS

The proposal of effective solutions to enable interoperability among heterogeneous and inter-organizational systems remains a key issue in the development of new Web-based applications, especially in the e-Business and e-Government contexts. The so called *Service-Oriented Architecture*, and more specifically its *Composition* layer, appears as a promising solution to deal with these demands. Even though most part of the academia and industry considers that *Dynamic* (and/or *Automatic*) *Service Composition* is the next stage to be reached, there is still a lot of misuse and misunderstanding regarding these concepts. In this paper, our first contribution is the proposal of a clear classification of different composition strategies with regards to their levels of dynamism and automatization.

Besides this classification, another important contribution of our paper is the introduction of a generic model driven approach for composing services, which is further specialized to illustrate the techniques that can be applied to enable fully dynamic service compositions. Aspects like dynamic selection of services and the use of AI techniques to automatically generate the execution plans are also discussed. An example service in the e-Government context is presented to illustrate the use of the proposed techniques. As already mentioned, the goal of this paper was not to focus on any specific technology or implementation solution, but rather to critically analyze the challenges and possible alternatives regarding service compositions in general. An interesting extension to this work would be to apply the proposed strategies in some technology specific scenario and evaluate its potentials and limitations considering different domains and platform characteristics.

We believe that much research is still necessary and also that a fully dynamic composition may be still far from becoming a reality, except for specific and well-defined application domains. Nevertheless, the research community, having recognized the potential of the evolving *Semantic Web*, has spawned several activities in the direction of *Semantic Web*

Services. With languages like the *Web Ontology Language* (OWL), machine-understandable Web service descriptions can be created and shared.

Generic service ontologies, such as *OWL-S* [4] and *WSMO* [8], in combination with appropriate rule languages and new AI techniques lay the foundations for semantically describing the functionality and the behavior of services, and keep the road open for a future of dynamic service compositions.

6. ACKNOWLEDGMENTS

The authors would like to thank the Brazilian Agencies CAPES and CNPq for financial support. Furthermore this work is also supported by the European Commission through the IST-1-002104-STP SATINE (Semantic-based Interoperability Infrastructure for Integrating Web Service Platforms to Peer-to-Peer-Networks) project.

7. REFERENCES

- [1] J. Bzivin, S. Hammoudi, D. Lopes, and F. Jouault. Applying mda approach for web service platform. In *Proc. of the 8th IEEE Intl Enterprise Distributed Object Computing Conference (EDOC 2004)*, pages 58–70, 2004.
- [2] F. Casati, S. Ilnicki, L.-J. Jin, V. Krishnamoorthy, and M.-C. Shan. eflow: A platform for developing and managing composite e-services. HP Labs Technical Report HPL-2000-36, HP Software Technology Laboratory, Palo Alto, CA, Mar. 2000.
- [3] CEFACT. *UN/CEFACT Modeling Methodology (UMM) User Guide CEFACT/TMG/N093*. http://www.unece.org/cefact/umm/umm_index.htm.
- [4] T. O. S. Coalition. Owl-s: Semantic markup for web services. White paper - <http://www.daml.org/services>, July 2004.
- [5] A. Dogac, Y. Kabak, G. Laleci, S. Sinir, A. Yildiz, S. Kirbas, and Y. Gurcan. Semantically enriched web services for the travel industry. *ACM Sigmod Record*, 33(3), Sep. 2004.
- [6] I. J. G. dos Santos and E. R. M. Madeira. Applying orchestration and choreography of web services on dynamic virtual marketplaces. *International Journal of Cooperative Information Systems (IJCIS)*, 15(1):57–85, Mar. 2006.
- [7] P. D. et al. *The Digital Ecosystem research Vision: 2010 and Beyond*. www.digital-ecosystems.org/events/2005.05/de_position_paper_vf.pdf, July 2005. Position Paper.
- [8] C. Feier and J. Domingue. *The Web Service Modeling Language WSML*. DERI International, WSML Final Draft, Apr. 2005.
- [9] M. Fluegge and D. Tourtchaninova. Ontology-derived activity components for composing travel web services. The Intl. Workshop on Semantic Web Technologies in Electronic Business (SWEB2004) - Berlin, Germany, Oct. 2004.
- [10] D. Frankel. Scaling the business process platform up. *MDA Journal* - <http://www.bptrends.com>, Dec. 2005.
- [11] IEEE. *IEEE Standard Computer Dictionary: A Compilation of IEEE Standard Computer Glossaries*. New York, NY, 1990.
- [12] ISO. *Use of UML for ODP system specifications*. ISO/IEC, JTC1/SC7 N3419. CD 19793.2, 2005.
- [13] M. Lin, H. Guo, and J. Yin. Goal description language for semantic web service automatic composition. In *Proc. of the 2005 Symposium on Applications and the Internet*, pages 190–196, 2005.
- [14] N. Milanovic and M. Malek. Current solutions for web service composition. *IEEE Internet Computing*, 8(6):51–59, Nov.-Dec. 2004.
- [15] OMG. *MDA Guide Version 1.01*. <http://www.omg.org/cgi-bin/apps/doc?omg/03-06-01.pdf>, 2003.
- [16] O. Patrascoiu. Mapping edoc to web services using yat1. In *Proc. of the 8th IEEE Intl. Enterprise Distributed Object Computing Conference (EDOC 2004)*, pages 286–297, 2004.
- [17] C. Peltz. Web services orchestration and choreography. *IEEE Computer*, 36(10):46–52, 2003.
- [18] S. R. Ponnekanti and A. Fox. Sword: A developer toolkit for web service composition. In *Proc. of the 11th WWW Conference, Honolulu, USA*, 2002.
- [19] J. Rao and X. Su. A survey of automated web service composition methods. In *Proc. of 1st Intl. Workshop on Semantic Web Services and Web Process Composition*, July 2004.
- [20] S. Ross-Talbot and N. Bharti. *Dancing with Web Services: W3C chair talks choreography*. <http://searchwebservices.techtarget.com/>, Mar 2005.
- [21] E. Sirin, J. Hendler, and B. Parsia. Semi-automatic composition of web services using semantic descriptions. In *Proc. of Web Services: Modeling, Architecture and Infrastructure workshop in conjunction with ICEIS2003*, 2003.
- [22] N. P. Tizzo, J. R. Borelli, M. de Jesus Mendes, L. Damasceno, A. Kamata, A. Figueiredo, M. A. Rodrigues, and J. G. S. Junior. Service composition applied to e-government. In *IFIP 18th World Computer Congress, Building the E-Service Society: E-Commerce, E-Business and E-Government*, pages 307–326. Kluwer Academic Publishers, 2004.