 OPAALS	OPAALS PROJECT Contract n° IST-034824
--	---

WP5: Integration with the Digital Ecosystem platform

Del5.7 - First increment of P2P infrastructure

 Information Society Technologies	Project funded by the European Community under the "Information Society Technology" Programme
--	---

Contract Number: IST-034824

Project Acronym: OPAALS

Deliverable N°: D5.7

Due date: 30/09/2009

Delivery Date: 30/09/2010

Short Description: This document is the implementation report relevant to series of deliverables (D3.1, D3.2, D3.3, D3.6 and D3.10) that aim to frame the infrastructural specifications in Digital Ecosystems research. As it has been shown (D3.10) in OPAALS the technological and social concerns in providing the necessary digital infrastructure are not treated as distinct but as part of the same continuum. This deliverable sets out the first implemented framework and clarifies the technical requirements for future work plans. This report introduces the Flypeer framework for transaction model in digital ecosystems. The framework tries to provide a fully distributed P2P environment, which executes different type of order service compositions. The proposed framework is considered at the deployment level of SOA, rather than the realisation level, and is targeted to business transactions between collaborating SMEs as it respects the loose-coupling of the underlying services.

Author: Amir Razavi, Fabio Serra and Paulo Siquira

Partners contributed: SURREY, IPTI

Made available to: ITA, LSE

Versioning		
Version	Date	Name, organization
0.1	05/01/2010	A. Razavi (SURREY)
0.2	10/01/2010	Fábio K. Serra and Paulo R C Siqueira (IPTI)
1.0	24/01/2010	A. Razavi and Paul Krause (SURREY), Fábio K. Serra and Paulo R C Siqueira (IPTI)
2.0	01/10/2010	A. Razavi (Surrey), P. Siqueira (IPTI)

Quality check

Internal Reviewers: Francisco José Lacueva Pérez, Jorge Veja-Murguía Merck

Dependencies:

Achievements*	<p>In this report we have presented a framework for executing distributed long-term transactions in Digital Business EcoSystems. Various forms of service composition have been considered in order to provide a closer representation of business transactions within a service-oriented architecture. Our model is considered at the deployment level of SOA while respecting the loose-coupling of the underlying services. Further, it addresses omitted results, through forward recovery, in a way that does not break local autonomy.</p> <p>Another strength of our approach is that it provides a light distributed coordinator which does not require to violate the autonomy of the local platforms or making any presumptions at the realisation level. This is particularly important in a business environment, especially for SME's. Further, these behaviour patterns can be applied dynamically on transactions in our approach, as was partly demonstrated in our example.</p>
Work Packages	<p>WP12: Task 12.8 – OS principles of communication and collaboration, Task 12.9, 12.10 – Business models in DEs</p> <p>WP11: Task 11.1 – collaboration and innovation in the Knowledge Economy, Task 11.4 – social innovation networks</p> <p>WP10: Task 10.10 – visualisation of P2P infrastructure</p> <p>WP3: Task 3.6, 3.7, 3.8 – Integration of Transaction model & P2P network services into the DE</p>
Partners	SURREY, IPTI, WIT
Domains	<p>Computer Science domain: P2P networks, interactions, long-running transactions, distributed systems and networks, redundancy, diversity, consistency, concurrency, design for failure, formal semantics, behaviour patterns, lock mechanisms, formal analysis, distributed identity, distributed trust.</p> <p>Social science domain: technology infrastructure for competitive advantage, lock-in, proprietary software / platforms, knowledge, openness, reciprocity, context-dependent trust, identity.</p> <p>Natural Science domain: simple reference to key analogies with ecosystems in nature and their key features found mostly in studies of biodiversity.</p>
Targets	Computer, social and natural science researchers, SMEs, business analysts. Computer science communities: database, transactions, P2P networking and applications, formal methods. Social science: language, socio-economics, governance, power and control, identity and trust.
Publications*	Razavi, A., S. Moschogiannis, and P. Krause. "An open digital environment to support business ecosystems." Peer-to-Peer

	<p>Networking and Applications, Springer New York. Available at: http://dx.doi.org/10.1007/s12083-009-0039-5.</p> <p>Razavi, A., P. Krause, and S. Moschoyiannis. 2010. "Digital Ecosystems: Challenges and Proposed Solutions." In Handbook of Research on P2P and Grid Systems for Service-Oriented Computing, IGI Global. http://www.igi-global.com/bookstore/Chapter.aspx?TitleId=40837</p> <p>Razavi A., Siqueira P., Serra F. and Krause P., 2010. "Flypeer: a JXTA implementation of DE transactions", 3rd International OPAALS (Open Philosophies for Associative Autopoietic Digital Ecosystems) Conference on Digital Ecosystems, 22-23 March 2010 Aracaju-SE-Brazil (In Press)</p>
PhD Students*	N/A
Outstanding features*	Despite the mentioned publications, the current theoretical framework has been proposed in several conferences and workshops from FESCA (at ETAPS) to IEEE-DEST conferences and we have proposed and presented several research tracks in these conferences.
Disciplinary domains of authors*	Computer Science

The information marked with an asterisk () is provided in order to address Recommendation n. 4 from the Year 2 review report*



This work is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 3.0 License. To view a copy of this license, visit : <http://creativecommons.org/licenses/by-nc-sa/3.0/> or send a letter to Creative Commons, 543 Howard Street, 5th Floor, San Francisco, California, 94105, USA.

TABLE OF CONTENTS

Executive Summary	6
Transactions in Digital Ecosystems.....	8
Transaction and Service Composition	10
Implementation Framework	12
Modeling & Executing a transaction	13
A simple scenario	14
Executing transaction	16
Comparison & Practicality.....	18
Conclusion and Further works	20
References.....	21

EXECUTIVE SUMMARY

The Digital Ecosystem is concerned with building an open source environment, through which businesses, in particular small to medium enterprises (SMEs), can interact within a reliable environment. The aim is to provide access to arbitrary services that help compose together to meet particular needs of the various partners. This collaborative software environment is being targeted primarily towards SMEs, who will be able to concatenate their offered services within service chains formulated on a digital ecosystem (Razavi, Krause, and Moschoyiannis 2006), (Razavi 2009).

Within a digital ecosystem a number of long-running multi-service transactions are expected to take place and therefore, of interest is the issue of providing support for long-term business transactions involving open communities of SMEs. A business transaction in this paradigm can be either a simple usage of a Web Service or a mixture of different levels of composition of several services from various service providers. We will argue that the current transaction and business coordination frameworks can lead to issues with tight coupling and violation of local autonomy for the participating SMEs.

Web Services are the primary example of the Service Oriented Computing (SOC) paradigm (Papazoglou 2003). The goal of SOC is to enable applications from different providers to be offered as services that can be used, composed, and coordinated in a loosely coupled manner. Web Services in fact provide a realisation of SOC. Although recent years have seen significant growth in the use of Web Services, there are some very significant technological constraints that are stopping their full potential from being realised within a digital ecosystem.

The conventional definition of a transaction (Date 2003) is based on ACID properties. However, as it has been indicated (Razavi, Moschoyiannis, and Krause 2007), in advanced applications these properties can present unacceptable limitations and reduce performance (Elmagarmid 1992).

The lack of consideration for the primary characteristics of SOC (such as loose-coupling) (Razavi, Moschoyiannis, and Krause 2007) or ignoring some important business requirements (such as partial results) (Razavi, Moschoyiannis, and Krause 2009), are important objections to the proposed transaction models, which also suffer from unnecessary complications of implementing a consistency model on top of the service-realisation boundary. In addition, the feasibility of a heavy coordinator framework causes some transaction models to use a centralised (or limited decentralised) coordination model (Furnis and Green 2005).

The JXTA technology is a set of open protocols that enable any device on the network to communicate and collaborate in a P2P fashion, creating thereby a virtual network. It provides the peers, on this network, to share and discover resources, services and exchange messages with other peers, even if some of these peers are behind firewalls and network address translations (NATs). Thus, JXTA is an important key, in the Flypeer implementation, in the sense of give an effort only in the Transaction Model implementation.

In the first section, we have provided an overview for transactions in Digital Ecosystems; section III has explained the role of service oriented infrastructure of digital ecosystems for transaction model. We introduce the Flypeer framework in section IV and section V provides the actual execution of transactions. Section VI compares the model with similar models and the last section provides conclusion and clarifies future work and roadmap for the project.

TRANSACTIONS IN DIGITAL ECOSYSTEMS

As we have already indicated, the very nature of business – as opposed to database – transactions, opens up a different angle from which to view transactions. For example, the specification of a transaction may involve a number of required services, from different providers, and allow it to be completed over a period of minutes, hours, or even days – hence, the term long-lived or long-running transaction. Indeed, a wide range of B2B transactions (business activities (Cabrera et al. 2004), (Razavi, Moschoyiannis, and Krause 2009)) have a long execution period. A business transaction between SMEs in a Digital Ecosystem can be either a simple usage of a Web Service or a mixture of different levels of composition of several services from various service providers.

It is important to stress that the long-term nature of execution frames the concept of a transaction in a digital ecosystem, since most usage scenarios involve long-running activities. In such cases, it is impractical, and in fact not feasible, to maintain full ACID properties throughout the lifetime of a long-running transaction. In particular, as we discussed in the Introduction, Atomicity and Isolation are questionable.

Within a digital ecosystem, a large number of distributed long-running transactions take place, each comprising an aggregation of activities which in turn involve the execution of the underlying service compositions. In such a highly dynamic environment there is an increased likelihood that some transaction (or one of its internal activities) will fail. This may be due to a variety of reasons (platform failure, service abort, temporary unavailability of a service, etc.) including the vulnerability of the network infrastructure itself (platform disconnection, traffic bottleneck, nodes joining or leaving the network, etc.).

The standard practice in the event of a failure is to trigger compensating actions that will effectively ‘undo’ the effects of the transaction – those effects visible before failure occurred. The objective is to bring the system to a state that is an acceptable approximation of the state it was before the transaction started. However, this is not a trivial task. Recovering the system in the event of a failure of a transaction (or an activity inside a transaction) needs to be done in a way that takes into account the dependencies both inside (to ensure all dependent activities are undone) but also outside the transaction (to ensure that any dependent activities of other transactions are also undone).

Further, and when considering SOA as the enabling technology for open e-business transactions, the recovery and compensation mechanism must respect the loosely-coupled nature of the connections, since interfering with the local state of the underlying service executions violates the primary requirement of SOA. In addition, access to the local state of another participants, may not even be possible in a business environment with SMEs as service providers. This is an issue that has been largely ignored by current implementations of transaction models such as Web Services Transactions (WS-Tx) (Cabrera, Copeland, Feingold, R. Freund, T. Freund, Joyce, et al. 2005), (Cabrera, Copeland, Feingold, R. Freund, T. Freund, Johnson, Joyce, Kaler, Klein, and Langworthy 2005) and Business Transaction Protocol (BTP) (Ceponkus et al. 2002).

The abortion of a transaction, even if it is successfully recovered and compensated for, can be very costly in the business environment. Rolling back the whole system in the event of a failure may lead to chains of compensating activities, which are time-consuming and impact on network traffic as well as deteriorate the performance. For this reason it is important to build into the system capability or flexibility to deal with failure. In other words, it is key to design for failure by adding diversity into the system and allowing for alternative scenarios. The idea is to get some leverage in avoiding the abortion of a whole transaction (and all other dependent transactions) by means of allowing alternative paths of execution in cases where the path chosen originally encountered a failure.

It is also important to be able to preserve as much progress-to-date as possible. If an activity (sub-transaction) of a transaction fails, it is essential to undo (roll back) only those activities that have used results of this activity, i.e. are dependent on it. It is highly desirable to avoid rollback of other activities that have produced results (committed) and are not dependent on the failed activity. These are often referred to as omitted results and do not need to be undone as that would mean they will need to be re-done (re-started, re-calculated, re-computed) once the transaction is restarted (after abortion and recovery). Addressing omitted results can have significant benefits for SMEs in digital ecosystems in terms of saving valuable time and resources.

In earlier work, we have performed an extensive review of Transaction models, such as WS-Tx and BTP, that have been designed with Web Services in mind, and are currently widely used in practice (Razavi and Krause 2009). Apart from certain issues regarding their coordination mechanism, which is geared towards centralised control, these models do not support partial results, do not provide capability for forward recovery, and there is no provision for covering omitted results.

TRANSACTION AND SERVICE COMPOSITION

Based on the specification advocated in (Yang, Papazoglou, and van den Heuvel 2002),(Singh and Huhns 2005),(Papazoglou et al. 2006) service composition can be considered along the following dimensions: data, process, security, protocol. In this paper we are concerned with providing P2P network support for distributed transactions and hence we will be concerned with the aspects of data and process composition. In general, security and protocol compositions are usually addressed on top of the transactional layer.

In particular, process-oriented service composition is concerned with the following aspects:

- Order: indicates whether the composition of services is serial or parallel.
- Dependency: indicates whether there is any data or function dependency among the composed services.
- Alternative service execution: indicates whether there is any alternative service in the service composition that can be invoked - alternative services can be tried in either a sequential or a parallel manner.
- Following (Yang, Papazoglou, and van den Heuvel 2002) these aspects can be seen within different types of service composition as follows.

Data-oriented service composition: The data generated at the service realisation level are released in terms of different data-objects. In this service composition, these data can be shared and manipulated between participants of a single transaction or, where partial results are concerned, be shared by participants of other transactions.

Sequential process-oriented service composition: This type of service composition invokes services sequentially. The execution of a component service is dependent on its previous service. These sequential dependencies can be based on commitment in which case we talk about Sequential with Commitment Dependency (SCD) where one cannot begin unless the previous service commits, or data in which case we talk about Sequential with Data Dependency (SDD) where one service relies on other service's outputs as its inputs.

Parallel process-oriented service composition: In this service composition, all the services can be executed in parallel. There may be data dependencies between them in which case we talk about Parallel with Data Dependency (PDD) or there may be differences in how and when the services can be committed (depending on some condition) in which case we talk about Parallel with Commit Dependency (PCD). When there are no dependencies between the parallel services we talk about Parallel without Dependency (PND).

Alternative service composition: This type of service composition indicates that there are alternative services to be deployed and one of them is necessary. They are categorised to two different types: Sequential Alternative Composition (SAc) where there is an ordering for deployment of these services, and Parallel Alternative Composition (PAc), where there is no

OPAALS Project (Contract n° 034824)

ordering (preference) between them and deployment of either service can satisfy the composition.

Generally, one or more service compositions can satisfy the user request. It can be seen that due to order and data in service composition, there can be increased complexity in composing services especially when transactions require a number of different services from different networked organisations. This means that there is a need for a context and data consistency model (at the deployment level) which can provide the correctness of the results.

IMPLEMENTATION FRAMEWORK

The Flypeer project brings a few things to the table: Web Services technologies, long running transaction support and peer-to-peer. To avoid having to re-inventing a lot of wheels, JXTA is used to provide the basis of the peer-to-peer architecture. With this in mind, the main goal of Flypeer is the supporting Long Running Transactions.

To achieve that, Flypeer abstract the communication details of JXTA from the service developer. Flypeer exposes an interface, *FlypeerBusinessService*, which the developer implements. This interface defines the features available to service developers and, more importantly, allow the framework to correctly plug it into the execution of a composition – more on that in the next section.

Deploying services in a peer-to-peer node, in Flypeer, is just a matter of creating a jar file with the proper service implementation classes and a configuration file, which follows the mechanism defined in the *ServiceLoader* class available in Java, version 6 and above.

The project also tries to be as close as possible to WS-* standards. This means, for example, that we use WSDL (actually a simplified version of it) for service description, and should support SOAP calls from the "external" world in a not distant future version. More than that, with each new version of the framework we are trying to get closer and closer to commonly used standards – allowing for new developers to use the project with a lighter learning curve.

MODELING & EXECUTING A TRANSACTION

In order to model a long running transaction, a service composition has to be created with the required parameters of its services and then the transaction can be executed. A class diagram of this process can be found in Figure 0.

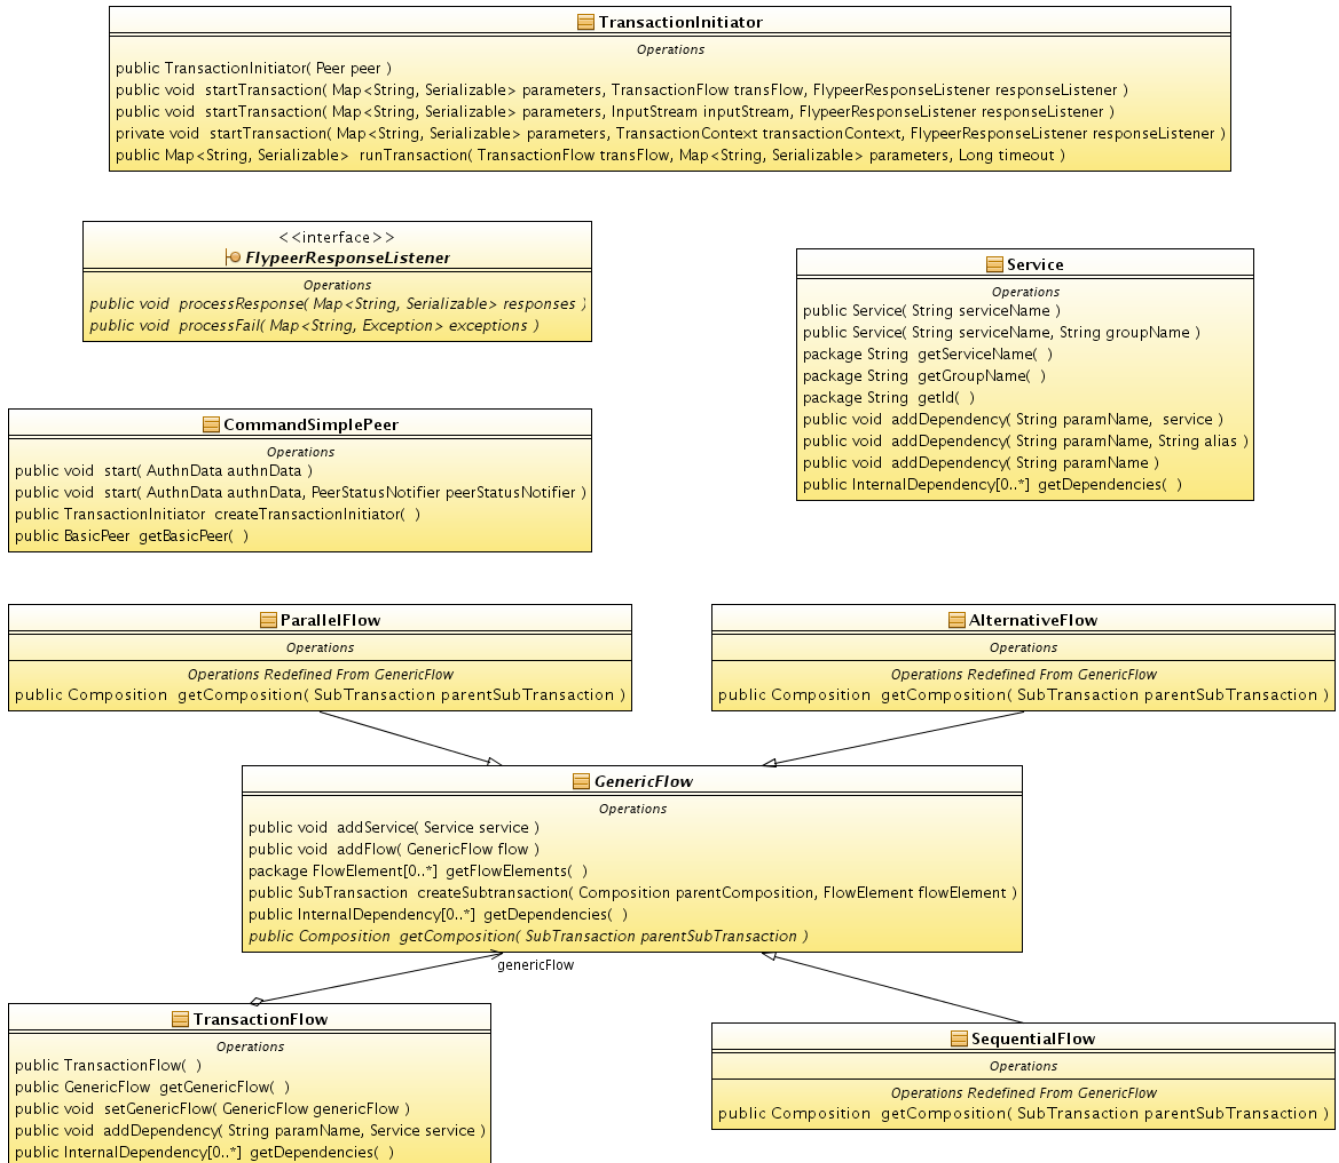


FIGURA 0 - CLASS DIAGRAM OF THE TRANSACTION EXECUTION

Regarding the service composition, Flypeer supports three types, SCD, PCD and Alternative Service compositions (recall section III). All of these can be mixed to create from simple to complex compositions, using an xml file or programmatically, in what we call Transaction Context.

The *TransactionFlow* is the class which represents the Transaction Context and where is created the transaction flow. A *GenericFlow* implementation must be added into it to begin the composition. This can be a *SequentialFlow*, a *ParallelFlow* or an *AlternativeFlow* object. A *GenericFlow* object can receive Service or/and another *GenericFlow* objects.

The *Service* class represents the service which will be called, and its dependencies. The dependencies are mapped using the *addDependency* method, passing the input name only, if it has no dependency with another service, or passing also the dependent service object that will receive the output of this service as its input.

After composing the transaction, all of the required service parameters must be added into the composition. This is done using a Map where the key is the input name and the value is the parameter value which has to be a *Serializable* object. These parameters are validated automatically, before the transaction is started, according to the input mapping described in the service WSDLs.

After that, the transaction can finally be started using a *TransactionInitiator* object, which is obtained through the *CommandSimplePeer* object. Then, all service coordinators of this composition are notified and receive the Transaction Flow and the parameters definition. The transaction starts with the initiator's coordinator calling the first service's one. After that, each service coordinator calls the next one, until the last service coordinator on the transaction is reached.

In the end of the transaction, the initiator receives the result in a callback object defined prior to transaction start. This called implements the *FlypeerResponseListener* interface.

One of the missing pieces of JXTA is the transaction support. This provides us a nice window of opportunity to contribute a new feature back to the community – which is likely to happen when Flypeer is fully implemented.

A SIMPLE SCENARIO

Fig 1 shows an example of a transaction in a tree structure with four services composed by three different composition types.




FIGURE 1 - TRANSACTION IN A TREE STRUCTURE

Fig 2 shows the transaction context (an xml file) for this example.

```

1  <?xml version="1.0" encoding="UTF-8"?>
2  <TransactionTree xmlns:xs="http://www.w3.org/2001/XMLSchema"
3      xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
4      transactionId="0001">
5      <SubTransaction subTransactionId="010">
6          <Composition compositionType="Parallel">
7              <SubTransaction subTransactionId="011">
8                  <Composition compositionType="Sequential">
9                      <SubTransaction subTransactionId="0111">
10                         <WebService serviceId="SERVICE1">
11                             <ServiceDescription>
12                                 Service 1
13                             </ServiceDescription>
14                         </WebService>
15                     </SubTransaction>
16                     <SubTransaction subTransactionId="0112">
17                         <WebService serviceId="SERVICE2">
18                             <ServiceDescription>
19                                 Service2
20                             </ServiceDescription>
21                         </WebService>
22                     </SubTransaction>
23                 </Composition>
24             </SubTransaction>
25             <SubTransaction subTransactionId="012">
26                 <Composition compositionType="Alternative">
27                     <SubTransaction subTransactionId="0121">
28                         <WebService serviceId="SERVICE3">
29                             <ServiceDescription>
30                                 Service3
31                             </ServiceDescription>
32                         </WebService>
33                     </SubTransaction>
34                     <SubTransaction subTransactionId="0122">
35                         <WebService serviceId="SERVICE4">
36                             <ServiceDescription>
37                                 Service4
38                             </ServiceDescription>
39                         </WebService>
40                     </SubTransaction>
41                 </Composition>
42             </SubTransaction>
43         </Composition>
44     </SubTransaction>
45     <Dependencies>
46         <InternalDependency paramName="param1" dependent="0111"/>
47         <InternalDependency paramName="param2" dependent="0112"/>
48         <InternalDependency paramName="param3" originator="0112" dependent="0121">
49         <InternalDependency paramName="param3" dependent="0122"/>
50     </Dependencies>
51 </TransactionTree>

```

FIGURE 2 - TRANSACTION CONTEXT

EXECUTING TRANSACTION

The transaction in Fig 1 can be modelled programmatically using the Flypeer Framework as it is demonstrated in Fig 3.


```

1 Service service1 = new Service("SERVICE1");
2 service1.addDependency("param1");
3 Service service2 = new Service("SERVICE2");
4 service2.addDependency("param2");
5 Service service3 = new Service("SERVICE3");
6 service3.addDependency("param3", service2);
7 Service service4 = new Service("SERVICE4");
8 service4.addDependency("param4");
9
10 SequentialFlow sFlow = new SequentialFlow();
11 sFlow.addService(service1);
12 sFlow.addService(service2);
13
14 AlternativeFlow altFlow = new AlternativeFlow();
15 altFlow.addService(service3);
16 altFlow.addService(service4);
17
18 ParallelFlow pFlow = new ParallelFlow();
19 pFlow.addFlow(sFlow);
20 pFlow.addFlow(altFlow);
21
22 TransactionFlow flow = new TransactionFlow();
23 flow.setGenericFlow(pFlow);
24
25 Map<String, Serializable> params = new HashMap<String, Serializable>();
26 params.put("param1", "1");
27 params.put("param2", "1");
28 params.put("param3", "0");
29 params.put("param4", "1");
30
31 commandSimplePeer = new CommandSimplePeer();
32 commandSimplePeer.start(getParameter("username"), new MyNotifier());
33 TransactionInitiator transInit = new TransactionInitiator(peer);
34 transInit.startTransaction(params, flow, new MyResponseListener());

```

FIGURE 3 - RUNNING THE TRANSACTION

First of all, there is a parallel composition which executes at the same time a sequential and an alternative composition.

The sequential composition runs first the Service1 and when it's finished the Service2 is called.

At this same time, the alternative composition also is being executed. First, the Service 3 is executed and if it runs with no errors, this composition is finished, but if something wrong happens the Service4 is called as an alternative for the Service3.

One thing to note about Service2 is that, as mapped in the Transaction Flow (Fig 2), the Service3 receives the output of the Service2 as an input parameter. So, in order to Service3 to be able to start its execution, Service2 has to have finished and sent its output to Service3 already.

When all compositions have finished the output is sent to the initiator as result of the transaction.

COMPARISON & PRACTICALITY

In Flypeer, the loosely coupled solution for a distributed transaction (Razavi, Moschoyiannis, and Krause), (Razavi 2009) has been described in the context of digital ecosystems, and in particular we have been concerned with services, transactions, and network support within the digital ecosystem initiative. The structure of the interaction network within the architecture we have proposed emerges through the local interactions that take place in the context of long-running business transactions. In contrast with our model, WS protocols use a customised coordinator for avoiding inconsistency.

For providing a consistent environment, during concurrent actions (service deployment and compositions), WS-* (WS-AtomicTransactions and WS-BusinessActivity) and BTP, are using the two-phase commit (2PC) protocol, which requires synchronisation for the phases. This is applied through a centralised coordination framework, based on WS-Coordination (Cabrera, Copeland, Feingold, R. Freund, T. Freund, Johnson, Joyce, Kaler, Klein, Langworthy, Little, et al. 2005). Fig. 4 shows a simple example of the use of WS-Coordination for executing a transaction where the Initiator creates a coordination context and the Participants, based on their registered services, deploy their respective services. The synchronisation for concurrency control is done in a centralised manner. This causes a single point of failure as well as a single dependency on the provider(s) of the centralised coordinator framework.

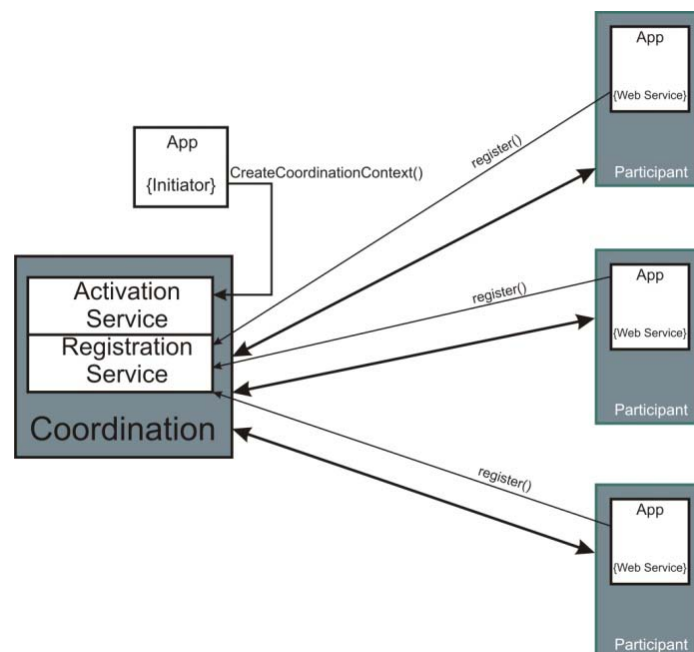


Figure 4 – WS-coordination

In addition, a more careful study of this coordination framework, such as that reported in (Furnis and Green 2005), shows it to suffer from some critical decisions about the internal build-up of the communicating parties - a view also supported in (Vogt et al. 2005). The

OPAALS Project (Contract n° 034824)

Coordinator and Initiator roles are tightly-coupled and the Participant contains both business and transaction logic. These presumptions are against the primary requirements of SOA, particularly loose-coupling of services and local autonomy, and thus are not suitable for a digital business ecosystem, especially when SMEs are involved. This is because smaller organisations tend to be more sensitive in revealing their local design and implementation precisely because this is often where their model lies (Razavi 2009), (Singh and Huhns 2005).

CONCLUSION AND FURTHER WORKS

In this paper we have presented a framework for executing distributed long-term transactions in Digital Business EcoSystems. Various forms of service composition have been considered in order to provide a closer representation of business transactions within a service-oriented architecture. Our model is considered at the deployment level of SOA while respecting the loose-coupling of the underlying services. Further, it addresses omitted results, through forward recovery, in a way that does not break local autonomy.

Another strength of our approach is that it provides a light distributed coordinator which does not require to violate the autonomy of the local platforms or making any presumptions at the realisation level. This is particularly important in a business environment, especially for SME's. Further, these behaviour patterns can be applied dynamically on transactions in our approach, as was partly demonstrated in our example.

The work on progress of Flypeer will support the previous locking mechanism of the digital ecosystems transaction model (purposed in D3.3, D3.10 and (Razavi, Moschoyiannis, and Krause 2009)). This is something we are keen to explore further as it lays the groundwork for a coordinated and collaborative service invocation specification to support long-term and multi-service transactions in Digital Ecosystems and will cover both topics of Internal Dependency Graph (IDG) and External Dependency Graph (EDG). The other part of Flypeer roadmap will provide a 'Trust framework' for businesses in digital ecosystems.

REFERENCES

- Cabrera, L., G. Copeland, M. Feingold, R. Freund, T. Freund, J. Johnson, S. Joyce, C. Kaler, J. Klein, and D. Langworthy. 2005. *Web Services Atomic Transaction (WS-AtomicTransaction)*. IBM, US: IBM. Available at: <http://www-128.ibm.com/developerworks/library/specification/ws-tx/#atom>.
- Cabrera, L., G. Copeland, M. Feingold, R. Freund, T. Freund, J. Johnson, S. Joyce, C. Kaler, J. Klein, D. Langworthy, M. Little, et al. 2005. "Web Services Coordination (WS-Coordination)." Available at: <http://download.boulder.ibm.com/ibmdl/pub/software/dw/specs/ws-tx/WS-Coordination.pdf>.
- Cabrera, L., G. Copeland, M. Feingold, R. Freund, T. Freund, S. Joyce, et al. 2005. *Web Services Business Activity Framework (WS-BusinessActivity)*. 2005. IBM DeveloperWorks.
- Cabrera, L. et al. 2004. "Coordinating Web Services Activities with WS-Coordination, WS-AtomicTransaction, and WS-BusinessActivity." Available at: <http://msdn.microsoft.com/en-us/library/ms996526.aspx>.
- Ceponkus, A. et al. 2002. "Business transaction protocol V1.0." *OASIS Committee Specification*, 3 June 2002.
- Date, C.J. 2003. *An Introduction to Data Base Systems*. 8th ed. Pearson Education.
- Elmagarmid, A. K. 1992. *Database transaction models for advanced applications*. Morgan Kaufmann.
- Furnis, P., and A. Green. 2005. "Choreology Ltd. Contribution to the OASIS WS-Tx Technical Committee relating to WS-Coordination, WS-AtomicTransaction, and WS-BusinessActivity (November 2005), lhttp." www.oasis-open.org/committees/download.php/15808.
- Papazoglou, M. P. 2003. "Service-Oriented Computing: Concepts, Characteristics and Directions." In *Proceedings of the Fourth International Conference on Web Information Systems Engineering*, Washington: IEEE Computer Society Press, December.
- Papazoglou, M. P. et al. 2006. "Service-oriented computing: A research roadmap." *Service Oriented Computing (SOC)*.

- Razavi, A., S. Moschoyiannis, and P. Krause. "An open digital environment to support business ecosystems." *Peer-to-Peer Networking and Applications*, Springer New York. Available at: <http://dx.doi.org/10.1007/s12083-009-0039-5> [Accessed September 9, 2009].
- Razavi, A. R. 2009. "DIGITAL ECOSYSTEMS, A Distributed Service Oriented Approach for Business Transactions."
- Razavi, A. R., and P. Krause. 2009. "Integrated autopoietic DE architecture - D3.10." Available at: http://files.opaals.org/OPAALS/Year_3_Deliverables/WP03/D3.10.pdf [Accessed January 15, 2010].
- Razavi, A. R., P. Krause, and S. Moschoyiannis. 2006. "Deliverable D24.5: DBE Distributed Transaction Model."
- Razavi, A. R., S. Moschoyiannis, and P. Krause. 2007. "A Coordination Model for Distributed Transactions in Digital Business EcoSystems." *Digital Ecosystems and Technologies (DEST 2007)*, IEEE Computer Society Press, Los Alamitos.
- Razavi, A. R., S. Moschoyiannis, and P. Krause. 2009. "An open digital environment to support business ecosystems." *Peer-to-Peer Networking and Applications Springer Journal*.
- Razavi, A. R., S. Moschoyiannis, and P. Krause. 2007. "Preliminary Architecture for Autopoietic P2P Network focusing on Hierarchical Super-Peers, Birth and Growth Models." Available at: http://files.opaals.org/OPAALS/Year_1_Deliverables/WP03/OPAALS_D3.1-final_submitted.pdf.
- Singh, M. P., and M. N. Huhns. 2005. *Service-Oriented Computing: Semantics, Processes, Agents*. Wiley.
- Vogt, F. H. et al. 2005. "Implementing Web Service Protocols in SOA: WS-Coordination and WS-BusinessActivity." *CECW* 5:21–28.
- Yang, J., M. P. Papazoglou, and W. J. van den Heuvel. 2002. "Tackling the Challenges of Service Composition in e-Marketplaces." In *Proceedings of the 12th International Workshop on Research Issues on Data Engineering: Engineering E-Commerce/E-Business Systems (RIDE-2EC 2002)*, San Jose, CA, USA.