ADAPT

IST-2001-37126

Middleware Technologies for Adaptive and Composable Distributed Components

Periodic Progress Report

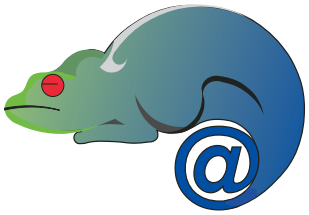

Deliverable Identifier:	PPR1
Delivery Date:	19th September 2003
Classification:	Restricted to the Commission and Appointed Reviewers
Authors:	Ricardo Jiménez-Peris, Marta Patiño-Martínez
Document version:	1.0, 18th September 2003
Contract Start Date:	1 September 2002
Duration:	36 months
Project Coordinator:	Universidad Politécnica de Madrid (Spain)
Partners:	Università di Bologna (Italy), ETH Zürich (Switzerland), McGill University (Canada), Università di Trieste (Italy), University of Newcastle (UK), Arjuna Technologies Ltd.(UK)

Project funded by the European Commission under the Information Society Technologies Programme of the 5^{th} Framework (1998-2002)

Contents

1 Executive Summary **3**

2 Project Rationale and Exploitation Plan **3**
2.1 Rationale . . . 3
2.2 Exploitation Plan . . . 3

3 Project Objectives **7**

4 Achievements and Project Status **7**
4.1 Basic Services . . . 8
4.2 Composite Services . . . 9
4.3 Evaluation Plan . . . 10
4.4 WS-CAF specification . . . 10

5 Adherence to Workplan **11**
5.1 Technical Progress . . . 11
5.2 Resource Usage . . . 11
5.3 Deliverable Schedule Update . . . 13
5.4 Modifications to Workplan . . . 14

6 Cooperation in the Project **14**
6.1 First plenary workshop. Madrid, 23-24th Sept. 2002 . . . 14
6.2 Second plenary workshop. Bologna, 13-14th Feb. 2003 . . . 14
6.3 Third plenary workshop. Zurich, 7-8th Jul. 2003 . . . 15
6.4 WP2 Meeting, Zurich, January 2003 . . . 16
6.5 WP1 Meeting, Bologna, 28 January 2003 . . . 16
6.6 WP2 Meeting, Newcastle, 22-23th May 2003 . . . 16

7 Coordination with other Projects/Programmes **17**

8 Promotion/Information Dissemination **18**
8.1 Standardization Bodies . . . 18
8.2 Publications . . . 18

9 Conclusion **19**

10 Appendices **20**
10.1 Updated Consortiums Relevant Persons Working and/or Associated to the Project . . . 20

1 Executive Summary

The ADAPT project is interested in developing support for the creation of adaptable web services. Web services have been proposed as a platform independent middleware solution that can interconnect components and applications across organizations. Two kinds of web services can be distinguished. Basic web services, those that do not rely on other services. And composite web services that invoke other web services to achieve their functionality. One of the main challenges being faced in the web service arena is how to achieve dynamically adaptable services. That is, services that adapt themselves to the changing environment. The way to achieve this adaptability is different in basic and composite services. For basic services, a generic infrastructure to build dynamic web services is needed. This generic infrastructure will enable the creation of web services that can adapt dynamically (i.e., while being online) to events such as site failures, site recovery, reconfigurations and changes in the load. On the other hand, adaptability in composite services takes a different form. Composite service adaptability consists in adapting the composition to changes in its constituent services. Additionally, what is especially interesting is the ability to predict the properties of the composition out of the properties of its constituent services.

2 Project Rationale and Exploitation Plan

2.1 Rationale

The presence of numerous basic services over the Internet is creating a new business opportunity for providing value added, inter-organisational services by composing multiple basic services into composite services. ADAPT will develop the technology, software infrastructure and a working system capable of defining, enacting, and monitoring inter-organisational business processes and supporting related coordination activities. The results of ADAPT will be open source and of interest to wide scientific and industrial communities, and will include: 1) Tools for creation of self-descriptive basic services built over commonly-used middleware; 2) a middleware platform for composite service creation and management, allowing enterprises to publish the services they provide, use services available globally in a secure manner as components of a composite service, specify its structure in a visual programming language, automatically compile the description into an executable process, and deploy and monitor the execution of such composite processes. The composite service middleware platform will contain re-configuration and self-repairing features to enable a composite service to adapt itself to changes caused by insertion or withdrawal of services, changes in network conditions and changes in user requirements. The project will develop demonstrator applications; the main demonstrator being a long running one (such as a travel planning e-service) through which the project will demonstrate composition, monitoring, management and dynamic adaptability to various changes both in the execution environment and in the system configuration.

2.2 Exploitation Plan

The exploitation of open-source software is usually carried out by one or more organizations that assume the maintenance of the product and provide support to its users. There are a number of companies that are devoted to provide consultancies on open-source software such as, Andago, Alcove, I+D Agora. We plan to publicize ADAPT results among them to foster their use. The clients of these companies are clients looking for open-source solutions and therefore they are very adequate for exploiting ADAPT results. Apart from communicating them the advances taken in ADAPT, we plan to invite them to the

open workshops organized by our consortium to guarantee an appropriate dissemination and foster the exploitation of ADAPT results by these organizations. In the Free/Open Source Software Projects Concertation Meeting organized by the Commission, we met a number of companies interested in exploiting ADAPT results. Just to name a few: IDEALX, Conecta, Jaluna, Minoru. These contacts will be very helpful for exploitation purposes.

A second group of organizations that can be interested in exploiting ADAPT results are those that maintain the software products that will be enhanced by ADAPT, such as PostgreSQL and JBoss. We think that an effective technology transfer is feasible in which the technology developed by ADAPT will influence future versions of their software. We will exploit the close contacts that the different partners keep with these organizations (McGill with PostgreSQL, Arjuna with JBoss).

Newcastle University.

Distributed Systems Group has a strong record of working with industries. In 1998, in conjunction with Nortel (Harlow research lab), we contributed to the development of the workflow standard by making a submission to the OMG based on our workflow technology. We collaborated with IBM, IONA Technologies and others in making an OMG submission Additional structuring mechanisms for the OTS; this submission has now been adopted as the OMG standard for extended transactions, and will be part of future release of J2EE middleware. We set up a company in 1998 in Newcastle to productise Arjuna transaction and workflow technologies. It is based in Newcastle, and began its life as privately owned, primarily by the five founding members (S. K. Shrivastava, and four research members of the Distributed Systems Group: S. J. Caughey, D. Ingham, M. C. Little, S. M. Wheater). The company was bought by Bluestone Software, Inc., a leader in business-to-Web and wireless technologies. In 2001, Hewlett-Packard bought Bluestone for $500M, and the original Arjuna company operated as HP-Arjuna Labs, a part of HP's middleware division. After a merger between HP and Compaq that led to substantial reorganisation of HP, Arjuna became independent again. Now based within the University campus, Arjuna Technologies (www.Arjuna.com) is a centre of excellence in transaction technologies and is focusing on building products to support reliable Web Services-based applications. Results from ADAPT will be used for existing and future research projects on middleware.

Arjuna Technologies Ltd.

Middleware is of importance to Arjuna Technologies as it is the foundation for building the new Internet generation of applications, and tying together disparate applications into a common solution for customers. Arjuna Technologies is a company that believes in open systems and customer choice. Arjuna Technologies has a partnership agreement with the JBoss group, an open-source J2EE provided, to supply transactioning technology. This partnership, in the future, may expand to the dissemination of ADAPT technology.

Arjuna Technologies views participation in the ADAPT project as a remarkable opportunity to gain considerable insights into what is achievable in the area of composition of electronic services. Results from ADAPT on basic services will help Arjuna Technologies enhance its products by supporting stateful services that are available and adaptable. ADAPT results on inter-organisational service composition will be of direct use to Arjuna Technologies work on future Web Services Business Integration products.

Arjuna Technologies has entered into an agreement with Hewlett-Packard by which technologies developed by Arjuna Technologies, can be exploited by Hewlett-Packard, and could be another channel for dissemination. The personnel of Arjuna Technologies have a tradition of working with universities

and industries for the development of open standards. Most recently, Arjuna Technologies personnel have worked with the research group at Newcastle (that is taking part in ADAPT), IBM, IONA and several others in the development of the OMG standard for extended transactions (activity service). It is taking that work further by contributing to the development of a similar standard for the J2EE platform (JSR 156 - Java API for XML Transactions).

In addition, Arjuna Technologies personnel have taken this expertise into the world of Web Services transactions and the OASIS Business Transactions Protocol (BTP) standard. BTP is specifically about loosening traditional ACID transaction semantics to facilitate the development of transactional e-commerce, leaving the semantics of what is meant to be transactional up to the Web Service and application to negotiate and understand. This was the first transaction standard for Web Services and attracted substantial interest from various areas of the development and business community.

However, the standards in this area are still to be agreed. Due to commercial pressures alternative standards to BTP have already been released by Microsoft and IBM and subsequently, because these standards are proprietary, Arjuna has been involved in a competing set of standards (Web Services Composite Application Framework) to push the evolution back within the open source arena. It is important, for the future of Arjuna Technologies, that a royalty free open standard for web service transactioning is widely adopted by the industry and the ability to author those specifications and create the necessary industrial momentum within the auspices of ADAPT has been greatly beneficial.

Università di Bologna.

The research group at the Department of Computer Science of the University of Bologna maintains close cooperation with national and international industries, including Microsoft (Cambridge Reasearch Laboratory) and Sun Microsystems, and national research institutes, the “Fondazione Marconi” and the ENEA (the Italian bureau for new technologies applied to the energy and environment), in the form of joint investigations under contracts and grant programmes. In addition, this research group is going to be involved in the national project entitled “Infrastructure Support for e-business applications”, that will be carried out in collaboration with a number of Italian Universities and companies. This project, which is currently under formal approval by the Italian Ministry of the University and Scientific and Technological Research, will greatly benefit from the results of the ADAPT project. In addition, two further projects, in which the research group in Bologna will be involved, can benefit from the ADAPT project results. Both these projects, entitled “A Distributed Broker for Quality of Service”, and “Middleware for advanced services over large-scale, wired-wireless distributed systems”, respectively, will deal with issues of QoS at the middleware level. These two projects have been recently submitted for approval to the Italian Ministry of the University and Scientific and Technological Research; both these projects will be carried out in collaboration with other Italian Universities and companies. The cooperation and collaboration channels mentioned above will be to transfer the results that will emerge from our research activity in the ADAPT project.

ETH Zurich.

The Information and Communication Systems Group at ETH Zurich has a long tradition of close interaction with industry and performing practically relevant research. Several systems developed by the group are now part of commercial products. For instance, results of past projects in the area of electronic commerce have been incorporated in IvyFrame, a business modeling and coordination tool developed by IvyTeam of Zug, Switzerland. Work on distributed transaction management resulted in a spin-off com-

pany that exploits the technology and the licensing rights of the related patents. Other systems developed by the group are public domain software. Work done on replication and high availability systems is being incorporated in PosgreSQL. Similarly, work on large data repositories performed in the context of the HESSI satellite of NASA is also publicly available and widely used in diverse scientific communities.

The results of ADAPT will be exploited in a similar vein. On the one hand, we plan to use the resulting system in normal graduate and undergraduate courses. On the other hand, the system will become an experimental test bed that will be used in cooperation with the local industry to evaluate technologies and build fast prototypes. This cooperation usually takes place in the form of consulting contracts where the research group used its own technology to demonstrate the feasibility of new concepts. Given the relevance of the topic, ADAPT will certainly be a very powerful platform for such projects.

Technical University of Madrid.

The Distributed Systems Lab (LSD) maintains a close cooperation with the national and international industry. One of the contributions of Lab to ADAPT, the database replication middleware, will be exploited in a wider context, thanks to a contract signed with Microsoft Research Cambridge that will fund the implementation of the infrastructure necessary to connect the database replication middleware with .NET and SQLServer. The Lab maintains contacts with national companies such as ELIOP and GridSystems. The cooperation with these companies has usually taken the form of contracts and/or preparation of joint research proposals. GridSystems is a world leader in developing support for grid systems. The results of ADAPT might be of interest to them to extend the scope of their product InnerGrid. ELIOP is a leader company in advanced control systems such as SCADA (Supervisory Control and Data Acquisition), Traffic Supervision in Highways and Tunnels, Railway and Underground Control Systems, etc. with clients in Latin America, North-Africa, Asia and Europe. Several contracts have been signed by our group with ELIOP during the last years. The results of ADAPT will be directly applicable to distributed control systems and SCADA systems. Additionally, the group has obtained regular funding since its creation in the form of research grants from the National Research Council dependent of the Science and Technology Ministry and from the Regional Research Council of Madrid (CAM). The most recently granted (2001-2004) national project funded by the Ministry is titled “Composable Fault Tolerance Services”. Two other Spanish universities collaborate in it. The project will greatly benefit from the results of ADAPT.

University of Trieste.

The Department of Electrical and Computer Engineering at the University of Trieste has a number of cooperations with local industries. The industrial environment of the North East of Italy is prevalently based on SMEs and is highly committed to R&D. Such an environment certainly embeds a great potential for diffusion of the results of ADAPT, in particular, due to the focus of the project on open-source products and inexpensive architectures based on COTS components. The Department, and specifically the group involved in ADAPT, has also strong cooperations with INSIEL, a leading software company in the design, development and maintenance of enterprise information systems for Local Authorities and the Health Service. The results of ADAPT will be applicable to a number of applications in this area. INSIEL is located in Trieste and belongs to the Finsiel Group, whose size and turnover make it the biggest software group in Italy. Finally, the group involved in ADAPT has just begun a new collaboration with Microsoft Research (Cambridge, UK) aimed at investigating new replication management strategies for highly-available databases. This channel will constitute an additional opportunity for diffusing the results of

ADAPT.

McGill University.

The distributed information systems group works together with members of the development team of the open source database systems PostgreSQL in order to integrate replication into PostgreSQL (see http://gborg.postgresql.org/project/pgreplication/projdisplay.php for the current open source effort). Furthermore, it has collaborations with several bioinformatics groups within McGill University (Department of Biomedical Engineering, Bioinformatics Lab of the School of Computer Science, Biochemistry Department) and Montreal (NRC Biotechnology Research Institute). All these partners are building information systems for managing their protein and experimental data. While access rights to the data are currently very restricted, future common projects will require more elaborate forms of interaction between the information sources of the different research groups and other research institutes within Canada. The results from ADAPT will be used to facilitate this cooperation.

3 Project Objectives

The main goal of ADAPT is to develop the technology and software infrastructure necessary for defining, enacting, and monitoring inter-enterprise business processes that are implemented as composite services with guarantees of availability, scalability and adaptability not only to changing network conditions and user requirements but also to reconfigurations and repairs. ADAPT will also provide middleware support for available and dynamically adaptive basic services that will be used to build higher level composite services.

4 Achievements and Project Status

The first year of the project has been devoted to:

1. Surveying the state of the art of standards and specifications (many of them being published during the first year of the project).
2. Familiarization with the software tools needed for the accomplishment of project objectives.
3. The design of the infrastructure for basic of composite service.
4. The development of languages for both specifying basic services and compositions.
5. Development of the evaluation plan.

These issues have been addressed in the first year deliverables. The first three items are dealt with in deliverables “D1 Basic Services Architecture”, “D5 Transaction Support”, and “D9 CS Middleware Architecture”. The fourth item has been addressed in deliverables “D6 Service Specification Language” and “D7 Composition Language”. Finally, the evaluation plan has been addressed in deliverables “D15 Updated Evaluation Plan” and “D16 Updated Evaluation Plan”.

4.1 Basic Services

The work during the first year in this part of the project was concentrated on surveying the state of the art, getting acquainted with the technology, and designing the architecture of the basic services infrastructure. In a first phase, the work was devoted to gaining experience with J2EE. During this phase the different available open source alternatives for the different tiers of J2EE were evaluated. This study led to the selection of different products, such as JBoss as the application server, Tomcat as the servlet engine, Axis as the SOAP engine, PostgreSQL as the database, and Spread as the group communication library.

A second phase of the work was concerned with the evaluation and study of the different products that were selected:

- JBoss.
- JBoss clustering.
- Comparison between the different group communication layers.
- The interaction between Tomcat and JBoss.
- The connection of JBoss with the database tier through JDBC.
- JBoss transactional support (JTA/JTS).

In a third phase various architectural options for the tiers involved were proposed and are being evaluated. Trieste has been developing JBora, a group communication toolkit that extends Spread and that should simplify the programming of replicated basic services. Mechanisms and policies for enabling JBora to tune some of its parameters automatically, i.e., for enabling JBora to adapt itself to unpredictable and varying run-time conditions, are being experimented. At the web tier level, Trieste has built a preliminary replication mechanism that is based on JBora and maintains state on behalf of remote clients. Trieste is also considering the feasibility of facilities for verifying on-line the consistency of such replicated state. At the application server tier, Bologna, McGill and Trieste have been considering several different options for its replication.

Finally, for the database tier two approaches have been explored. One consisted in the replication within the database, and the other in the replication at the middleware level. For the replication within the database, a new version of Postgres-R is being produced that works with PostgreSQL 7.2 and provides master-slave replication. The version is at a prototype stage, but it is being tested worldwide through the Postgres open-source community.

On the other hand, the middleware approach has been concerned with the exploration of dynamic adaptability aspects at the database tier and its connection to JDBC. In particular, in addition to tolerate failures, it is being studied how to perform dynamic load balancing (to adapt to changes in the load), adaptive control admission (to adapt to changes in the workload), and adaptive online recovery (to perform recovery without stopping servicing requests and devoting more or less resources to recovery depending on the current load). The latter issue is the most advanced of the three and the algorithm has already been written and formally proven in a paper. The support for the three attributes is currently under development. With regard to JDBC connectivity, a prototype of the middleware with support for JDBC clients is close to completion, for autocommit transactions. Upon completion of JDBC support, it will be necessary to develop support for non-autocommit transactions. This feature will be crucial to enable the integration of the work performed on the application server and database tiers and to achieve full generic support in the database tier.

4.2 Composite Services

Similarly to basic services, the work during the first year in this part of the project was concentrated on surveying the state of the art, getting acquainted with the technology, and designing the architecture of the composite services infrastructure.

In the first phase, the work the state of the art in standards, specifications and available tools were surveyed. The surveyed standards and specifications have evolved continuously since the proposal was accepted. Among the studied specifications are:

1. Web service description language (WSDL) for specifying web services.
2. SOAP as service invocation standard.
3. Composition languages such as Web Service Flow Language (WSFL), Microsofts XLANG, Business Process Modeling Language (BPML), Business Process Language for Web Services (BPEL4WS), and ebXML.
4. Conversation and choreography specification languages such as Web Service Conversation Language (WSCL) and Web Service Coreography Language (WSCI).
5. Transaction coordination specifications such as Business Transaction Protocol (BTP), WS-Coordination/WS-Transaction (WS-C/WS-T) and WS Coordination Activity Framework (WS-CAF).
6. Publishing specification such as Universal Description, Discovery and Integration (UDDI).

Regarding the studied tools, the main ones were Tomcat as web server engine, Axis for deploying and invoking web services, JOTM and JBoss transaction managers as JTA/JTS compliant transactional engines, a UDDI registry implemented in Java and BPWS4J, the IBM interpreter of BPEL4WS. In order to start experimentation with inter-organizational services, a simple web service was deployed at different partner sites so it could be invoked by a simple composite service.

In a second phase, we were concerned with answering some questions: a) how are composite services created; b) which types of composite services can be created; c) how are composite services enacted. The conclusion for question a) was that composite services can be created either top-down or bottom-up. In a top-down approach, the composite service is first created and then, it is associated by means of invocations to available web services. Whilst in the bottom-up approach, first the constituent web services are identified and selected by means of a browser, and then they are linked together by the business logic of the composition. It was decided to provide support for both styles of composition.

Concerning question b), two kinds of composite services were identified: 1) A composite service is run by a single organization that might invoke constituent services at other organizations; 2) A multi-party composite service that is run cooperatively by a set of organizations.

Finally, with reference to the enactment of composite services, two types were identified: 1) a centralized approach in which a composite service is run entirely at a single site; 2) a distributed approach in which a composite service is run in a distributed fashion across a set of sites. For the first approach, the workflow engine from ETHZ (extended with web service capabilities) will be used, whilst for the second approach, UNEW is developing a Java-based engine.

The third phase was focused on the design of the basic infrastructure for composite services. This basic infrastructure consists of the following subsystems: 1) A centralized composite service engine; 2) A distributed composite service engine; 3) A transactional engine for transactions across organizations; 4) A visual composition tool that would act as development environment for composite services.

Finally, in a fourth phase our efforts were devoted to develop a language for the description of web services and a language for the composition of web services. For the description of web services, WSDL was selected to its widespread acceptance. It has been extended through its extensibility elements to incorporate a rich set of attributes related to QoS attributes such as sequencing constraints (conversations), transactional properties, and performance metrics.

4.3 Evaluation Plan

The "Updated Evaluation Plan includes an analysis of the existing and proposed standards and technologies to identify the technical capabilities and attributes, which are relevant to the goals of the ADAPT project. These goals have been identified as: availability, scalability, performance, service description, composability, configuration, adaptation, process definition, replication, security, transaction models.

The set of selected existing and proposed standards and technologies that have been identified are: Java 2 Platform, Enterprise Edition (J2EE), Business Transaction Protocol (BTP), CORBA Component Model (CCM), Web Services (WS-Transaction) and Web Services (WS-Coordination), Business Process Execution Language for Web Service (BPEL4WS),Web Service Choreography Interface (WSCI),Web Service Flow Language (WSFL), Microsofts XLANG,Web Service Conversation Language (WSCL), Web Service Description Language (WSDL), WS-Security (WSSec), Universal Description, Discovery and Integration (UDDI), OMG Interface Definition Language (OMG IDL), CORBA Interface Repository (CORBA IR), and the Microsoft Cluster Service (MSCS).

The evaluation plan will address the functional capabilities and attributes of each of these technologies and derive a set of questions that can be used to evaluate the functionality of ADAPT.

4.4 WS-CAF specification

The Web Services Composite Application Framework specification was released on 28 July 2003. WS-CAF is a collection of three specifications – Web Service Context (WS-CTX), Web Service Coordination Framework (WS-CF), and Web Service Transaction Management (WS-TXM) designed to solve problems that arise when multiple Web services are used in combination (composite applications) to support information sharing and transaction processing.

Currently, the BTP standard has lost traction due to the release, for commercial reasons, of the Microsoft and IBM backed Web Services Coordination and Transactions standards. WS-C/T remains proprietary so it is essential that the industry has a royalty free alternative. The release of the Arjuna and IONA authored WS-CAF specifications has achieved this. The task now is to open up the participation and evolution of a single standard for transaction coordination to other companies (including MS and IBM) to ensure a non-proprietary (open source) standard for the industry. Industry analysts have already looked favourably upon this initiative - see:

`http://www.cbdiforum.com/public/news/index.php3?id=1288&news_date=`

An extract from the above referenced article is included below.

CBDI Commentary.

Last August IBM, Microsoft and BEA published two specifications WS-Transaction and WS-Coordination. Now we have a further specification covering a very similar area from Oracle, Sun and partners. At a capability level these specifications cover very similar functionality.

At first sight the WS-Coordination specification looks to be tightly bound with BPEL, which of course Oracle and Sun are currently in opposition to. Whilst we are predicting that the broader industry will eventually converge around a BPEL like protocol, it is clear that this is currently very inadequate when compared to ebXML, and it will not happen soon. However in reality the interface between business process steps is a Web Service, and hence an agreed standard protocol.

Neither specification has been submitted to a standards body yet. And this latest publication from Oracle Sun et al. seems like an attempt to bounce IBM et al. into a standards process that has broader representation. We think this is a good idea. The industry really needs to address how it is going to come together on the standards in the area of complex business transactions, and the arbitrary leadership of IBM, Microsoft plus a chosen partner (in this case BEA) is really inadequate for what is a complex and crucial set of protocols. We assess that the specifications are sufficiently close to be a basis for agreement on that specific area, without requiring the more difficult and contentious area of process scripting to be resolved concurrently.

5 Adherence to Workplan

5.1 Technical Progress

The workplan is progressing according schedule.

5.2 Resource Usage

- UPM.
  - Staff: 31.82 person-months (scheduled 33.3 person-months).
  - Resource usage: 86,905 Euros (scheduled 108,722 Euros).

  During the first year, due to a problem to contract an engineer as foreseen, two part-time master students were contracted instead. For the second and third year an engineer is already being contracted (plus one part-time master student). The lack of technical qualification during the first year was in part compensated by an increased participation of the permanent staff.

- UniBo.
  - Hired staff: 20 person-months (scheduled 24 person-months).
  - Permanent staff: 16 person-months (scheduled 12 person-months).
  - Resource usage: 68,069 Euros (scheduled 81,810 Euros).

  There was a 4-month delay in the start date of the second scientist (Jaksa Vuckovic) due to administrative delays in the hiring process. This shortage has been partially compensated by increased effort by the permanent staff during the period 1/9/2002 - 31/12/2002. There has been no modification to the scheduled workplan or the delivery dates.

- ETHZ.
  - Hired staff: 18 person-months.
  - Resource usage: 75,034 Euros (with no overheads).

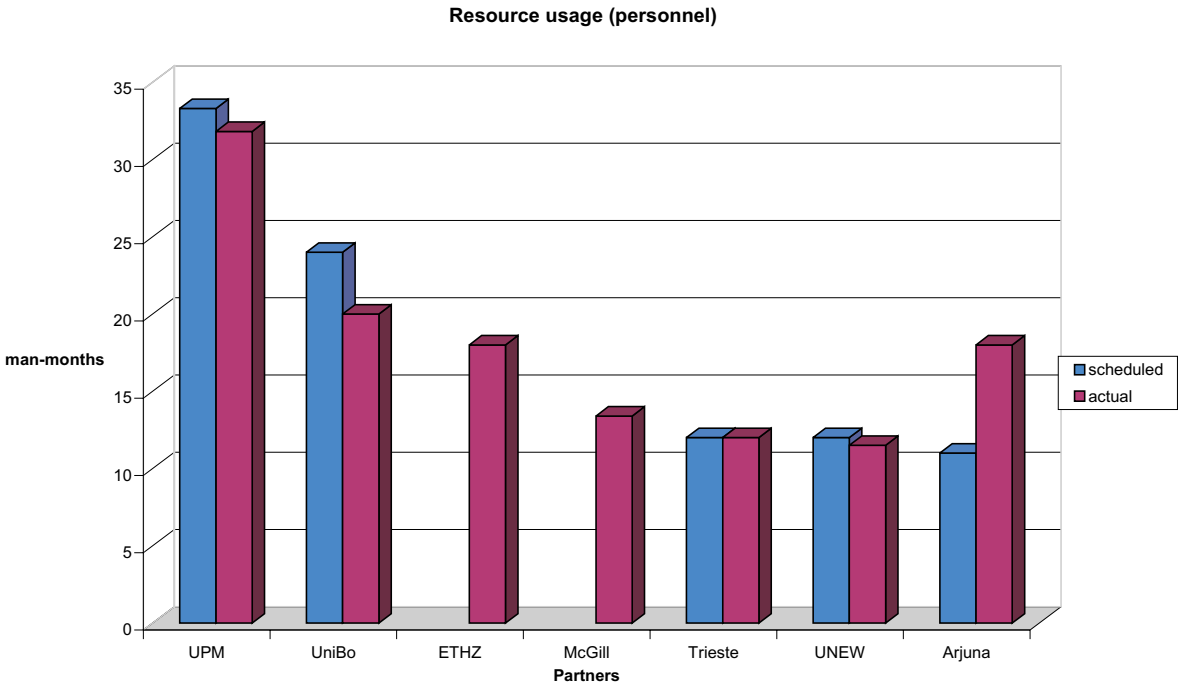

Figure 1: Resource usage (personnel)

- McGill.
  - Hired staff: 12 months by a PhD student, 1.4 months by a technician.
  - Permanent staff: 1.2 months.
  - Resource usage: 32,248 Euros (with no overheads).

  McGill receives partial funding from the Quebec government for this project. However, the money came in late. Hence, part of the work to be performed by the technician (1.4 months) within the first year will be performed in the second year.

- Trieste.
  - Hired Staff: 12 person-months (as scheduled).
  - Permanent Staff: 4 person-months.
  - Funding usage: 42,809 Euro (scheduled 42,878 Euro).

  There was a 3-month delay in the hiring of Milan Prica (a Croatian citizen), due to the intricacies of local rules for the hiring process of foreign people. This delay has been compensated by increased effort of the permanent staff and by short-term contracts to other people.

- UNEW.

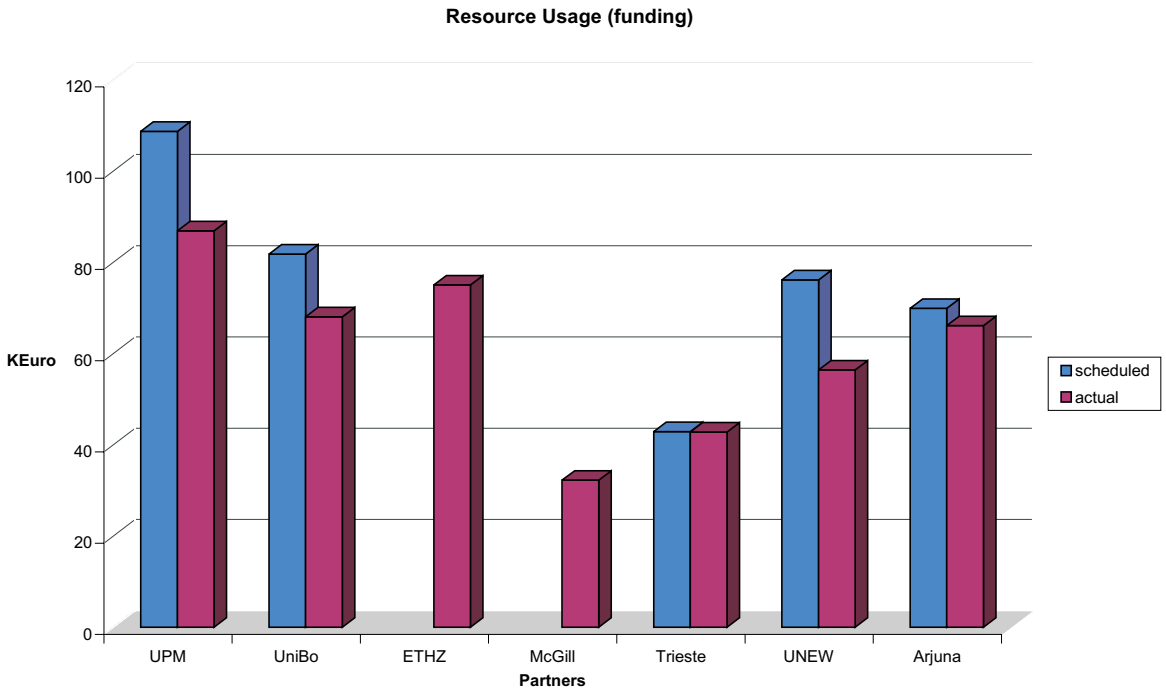

Figure 2: Resource usage (funding)

  - Hired Staff: 11.5 person-months (scheduled 12 person-months).
  - Permanent Staff: 12 person-months.
  - Funding usage: 56,408 Euro (scheduled 76,142 Euro).

- Arjuna.

  - Staff: 18 person-months (scheduled 11 person-months).
  - Funding usage: 66,118 Euro (scheduled 69,904 Euro)

  ATL faced significant administration difficulties over the first 6 months of this project negotiating an amendment to the contract with HP and the Commission to become a contractor on ADAPT (in place of HP) and subsequently carried out only 3 man-month work during the first 6 month period. During the second 6 month period we have been able to allocate extra resource on ADAPT to ensure all deliverable obligations were met and to push forward work on the WS-CAF specification. Despite contributing extra effort to ADAPT within the first 12 months our cost statement remains well within the original budget allocation.

5.3 Deliverable Schedule Update

The deliverable schedule has not been changed. One of the month 6 deliverables, “D5 Advanced Transaction Model” has changed its name to better reflect its contents. After the study of the state of the art

of advanced transaction models it became clear that no new transaction model was needed and that no single advanced transaction model can suit any arbitrary CS. The decision was to provide flexible support for a wide range of advanced transaction models. Therefore, we decided to extend the scope of the deliverable to cover: the state of the art of transactional support for BSs and CSs (including advanced transaction models), standards for transaction support, and architecture outline of transactional support for BSs and CSs.

5.4 Modifications to Workplan

No modification has been made.

6 Cooperation in the Project

6.1 First plenary workshop. Madrid, 23-24th Sept. 2002

The first plenary meeting was devoted to plan the first stages of the project. During the first day of the meeting the partners made presentations on their respective backgrounds related to the project. The partners also pointed out which could be their contributions to the project.

During the second half of the day, two issues were dealt with. The first discussed issue was how to proceed with the industrial partner substitution. Three possibilities were considered: to continue with a different lab of HP, to substitute HP by a new company formed out of the staff of the HP Arjuna Lab, or to search for another partner. The option considered more plausible was to substitute HP by the new company funded by the staff of the former HP Arjuna Labs (ATL). This option was the less disruptive as it meant to continue with the same team of people.

The rest of the day and the next half day were devoted to set up a detailed workplan and to assign responsibilities regarding the month 6 deliverables. Additionally, the main pieces of software to base the work of the consortium were decided: JBoss as application server, Tomcat as servlet engine and Axis as SOAP engine.

Talks and schedule are available at: http://adapt.ls.fi.upm.es/madrid-sep-2002-meeting.htm

6.2 Second plenary workshop. Bologna, 13-14th Feb. 2003

This second plenary workshop was mainly devoted to come out with a common and integrated vision of the project, as well as to discuss the main architectural decisions taken during the first six months of the project. The meeting was extremely fruitful and helped to solve issues that could not be solved by e-mail.

During the first day, it was discussed the architecture of the basic services (BSs), especially the main issues regarding the interface between the different components. During the first half of the second day, the architecture of the composite services was considered. One of the main issues discussed was whether to pursue a centralized or decentralized (workflow) engine for supporting CSs. The decision was to continue, at least initially, with both approaches to see which one was more promising. During the second half of the second day, issues related to the integration of BSs with CSs were discussed.

Talks and schedule available at: http://adapt.ls.fi.upm.es/bologna-feb-2003-meeting.htm

6.3 Third plenary workshop. Zurich, 7-8th Jul. 2003

The workshop was mainly devoted to discuss the integration about the different lines of work carried out so far. For this purpose, the whole workshop was organized as a series of round tables instead of the previous run of talks organization. The discussion about the integration was performed in a bottom-up fashion. It was first discussed the integration within basic and composite service levels, and then it was discussed the integration between basic and composite services.

In order to center the discussion, the first round table was devoted to revisit the overall objectives of the project. Then, the rest of the day was devoted to discuss the objectives and integration of both basic and composite services. The discussion started first with composite services as a follow-up of the WP2 meeting held in Newcastle. The conclusions of this meeting were summarized for people not attending that meeting. The different developments achieved by each of the partners regarding composite services were discussed such as: the prerelease of the visual composition tool developed by UPM, the study of the specification of real compositions from xCBL (XML Common Business Library) as conversations performed by ETHZ, the study of the specification of sequencing constraints with PI-Calculus performed by UNEW, the study of the specification of transactional attributes performed by UPM, and the composition language being developed by UNEW.

One of the conclusions of the discussion was that the service specification language deliverable would be complete enough discussing both sequencing constraints and transactional attributes. It was also acknowledged that performance attributes such as response time or throughput, although interesting they have a dynamic nature and are not adequate for a static specification of a service. However, it might be possible to include links to web services that can be used to query about a specific attribute.

Another conclusion from the discussion was the need for providing support in the visual composition tool for building compositions bottom-up. That is, that it might be possible to choose services from a repository, such as UDDI, and then, glue them into a composition. Another comment regarding the prerelease of the visual tool was that it was the current complete view was too crowded. Therefore, it was necessary to provide different views, each projecting particular elements of the composition. Additionally, it was also acknowledged that a more abstract view (e.g. showing message interactions and basic structuring such as sequential/parallel composition) would also be more adequate as visual representation of the composition language developed by UNEW.

Finally, it was also concluded that WSCL was insufficient to specify complex conversations. More concretely, it was identified a basic problem that was the one of specifying conversations with asynchronous interaction for which the existing languages supporting conversation specification were not adequate.

In the basic services session, the main topics debated were: the problems of achieving exactly-once semantics in the application server and the integration between the replicated application server with the replicated database layer. In the discussion of the exactly-once semantics, the suite of algorithms designed by Bolonia, McGill and Trieste were presented. The basic premise of the algorithms is that each client request was executed under the umbrella of an atomic transaction. However, a problem was identified regarding the interaction of the transaction manager of the application server and the database in the advent of failures. The conclusion was that this problem could be solved by making the transaction manager of the application server, replication aware.

The recent developments in the undergoing implementation of the database replication middleware by UPM were also presented. Then, it was discussed the integration between the replicated application server and the database replication middleware. It was found out that the current approaches could not be integrated, since the middleware only works in auto-commit mode and the replicated application

server requires a transaction enclosing the full interaction with the database. The conclusion was that by setting the application server to container managed persistence, the middleware could then be extended to deal with transactions spanning several messages as required by the replication policy in the application server.

The next day, it was discussed the integration between basic and composite services. One of the central issues in the discussion was the link between transactional composite services and basic services. It was concluded that the transactional engine to be developed by UPM would provide the glue to coordinate transactions between a composite service and the local basic service participating in it. Another conclusion drawn from the discussion was that the link between the basic of composite services was that the basic services partners would provide an infrastructure to create basic services with a rich set of QoS attributes such as high-availability and transactional semantics that could be exploited successfully exploited by a composite service making use of them.

Another session of the day was devoted to plan the review meeting. In particular, the draft schedule of the meeting was produced to be sent for approval to the project officer.

The executive board also met during the workshop. One of the decisions taken by the board was the need of holding an early plenary workshop to prepare for the deliverables of month 12. It was decided to hold it in Bologna in middle December.

Schedule available at: http://adapt.ls.fi.upm.es/zurich-jul-2003-meeting.htm

6.4 WP2 Meeting, Zurich, January 2003

The meeting was devoted to discuss the possible alternatives for the architecture of the CS services. The results of this meeting were taken as input for the deliverable “D9 CS Middleware Architecture”.

6.5 WP1 Meeting, Bologna, 28 January 2003

The technical program of the meeting consisted of two parts: discussion of J2EE and JBoss, and of group communication infrastructures in general.

The main issues discussed at the meeting were strategic: specifically, how to structure WP1 so that all partners can work together and individually as developers and researchers.

6.6 WP2 Meeting, Newcastle, 22-23th May 2003

The meeting was devoted to the understanding of how composite services should be specified and built. In particular, the participants were interested in a case study in which a particular business interaction is automated by specifying a composite service among different partners. Then, each partner should develop the code of its role fulfilling the global interaction specification. The participants were interested in devising and developing some kind of specification from which it would be possible to derive skeleton business processes with the corresponding part of the protocol hardwired into the business process definition. Then, each partner only had to add the business logic to the automatically derived process definition with the guarantee that the specification of the global process would be fulfilled. The participants decided to study the different languages available for specifying conversations and multi-party interactions and experiment with real examples such as those part of xCBL (XML Common Business Library).

7 Coordination with other Projects/Programmes

Different coordination activities with other projects have been undertaken up to the end of the period:

- Free/Open Source Software Projects Concertation Meeting, Brussels, 24th June 2002. A representative of the coordinator, prof. Marta Patiño-Martínez, attended the meeting and made a presentation of the ADAPT project and established contacts with some industry representatives interested in the potential results of the project.
- IST Cluster Meeting for Service Engineering Functionality, Bucharest, Oct. 11th 2002 in conjunction with the "Information Society Technologies for Broadband Europe". The coordinator, prof. Ricardo Jiménez-Peris attended the meeting and made a presentation of the ADAPT project.
- IST Cluster Meeting for Service Engineering Functionality, Brussels, 17th February 2003. The coordinator, prof. Ricardo Jiménez-Peris, attended the meeting. This meeting followed up the one held in Bucharest. The meeting resulted in the agreement of preparing a proposal for a coordinating activity for the first call of FP6. This proposal will study the different perspectives of the concept of a service and how services are engineered in different contexts. The partners of ADAPT will participate in the proposal.
- MIDAS Workshop on "Middleware for Composable and adaptable Services", Montreux, 14-15 November 2002. The workshop was attended by three partners of the ADAPT consortium (prof. Santosh Shrivastava from UNEW, prof. Ricardo Jiménez-Peris from UPM, prof. Gustavo Alonso from ETHZ). The MIDAS consortium is led by prof. Santosh Shrivastava what guarantees a fluent communication between the two projects. The invited ADAPT partners participated actively in the meeting. More concretely, prof. Gustavo Alonso presented an invited talk and all the ADAPT partners participated in the working groups organized as part of the meeting.
- DeFINE Workshop on Dependability Foundations for Information infrastructures - Network of Excellence, Pisa 25-27 November 2002, Paris 24th of February 2003, and London 2nd of April 2003. Two partners of ADAPT were invited to the meeting (UPM, ETHZ) and were finally selected to be part of the membership of a NoE proposal for the first call of FP6 on the dependability objective.

8 Promotion/Information Dissemination

8.1 Standardization Bodies

Arjuna Technologies is in the process of joining the following standards bodies:

- World Wide Web Consortium (W3C): "The World Wide Web Consortium (W3C) develops interoperable technologies (specifications, guidelines, software, and tools) to lead the Web to its full potential. "
- OASIS: "OASIS is a not-for-profit, global consortium that drives the development, convergence and adoption of e-business standards."
- Java Community Process (JCP): "Java Community Process (JCP) is the way the Java platform evolves. It is an open organization of international Java developers and licensees whose charter is to develop and revise Java technology specifications, reference implementations, and technology compatibility kits."

These standards bodies are active in areas in which ADAPT will be contributing, such as service and process definition. Arjuna Technologies intends to champion the technologies produced by ADAPT in these standards bodies. As aforementioned Arjuna has been a lead participant in the WS-Caf specifications for coordinated web services.

8.2 Publications

The project publications are the following:

- G. Alonso. *Myths around Web Services*. Bulletin of the Technical Committee on Data Engineering. IEEE Computer Society. pp. 3-9. Dec. 2002.
- R. Jiménez-Peris, M. Patiño-Martínez, G. Alonso. *An Algorithm for Non-Intrusive, Parallel Recovery of Replicated Data and its Correctness*. 21st IEEE Int. Conf. on Reliable Distributed Systems (SRDS 2002), pp. 150-159. Oct. 2002, Osaka, Japan.
- A. I. Kistijantoro, G. Morgan, S. K. Shrivastava and M.C. Little, *Component Replication in Distributed Systems: a Case study using Enterprise Java Beans*, 22nd IEEE/IFIP Symp. on Reliable Distributed Systems (SRDS2003),**?**Florence, Oct.**?**2003.
- S. J. Woodman, D. J. Palmer, S. K. Shrivastava, S. M. Wheater. *A System for Distributed Enactment of Composite Web Services*, Work in progress report. Int. Conf. on Service Oriented Computing, 2003.
- A. Bartoli, B. Kemme, V. Maverick, A. Montresor, J. Vuckovic, and H. Wu. *Replication techniques for J2EE*. Technical Report. July 2003.
- G. Masarin, A. Bartoli and V. Maverick. *On-line consistency checking for replicated objects*. Int. Conf. on Distributed Objects and Applications (DOA), 2003.
- A. Bartoli, C. Calabrese, M. Prica, E. Antoniutti di Muro, A. Montresor. *Adaptive message packing for group communication systems*. Workshop on Reliable and Secure Middleware, 2003.

Some related publications are:

- W. Bausch, C. Pautasso, G. Alonso. *Programming for Dependability in a Service-based Grid*. The 3rd IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid 2003). May, 2003, Tokyo, Japan.
- R. Jiménez-Peris, M. Patiño-Martínez, G. Alonso, and B. Kemme. *Are Quorums an Alternative to Data Replication?* ACM Transactions on Database Systems, Vol. 28, N. 3, Sept. 2003, pp. 257-294, ACM Press.

Other dissemination activities include:

- The project was presented to representatives from Microsoft Research Cambridge during a visit to UPM in Autumn. As a result, Microsoft became interested in the database replication middleware, what resulted in signing a contract in Summer 2003 to adapt it to .NET and SQLServer.
- The project was also presented to industry in a two-day course imparted in Zurich by ETHZ (10-11 February, 2003) with 41 participants from diverse industrial sectors.
- Publish the project in the LEARN web page, part of an accompanying measure in the IST programme to inform people about European research into information technology and to encourage young people to consider science and technology careers.
- Talk presenting the project and related activities at ObjectWeb (INRIA Grenoble) by Prof. Gustavo Alonso. March, 2003.
- Prof. Gustavo Alonso participation in a panel on Web Services at VLDB 2003 (Berlin, Sept. 2003)
- Book on Web Services co-authored by Prof. Gustavo Alonso (”Web Services: Concepts, Architecture and Applications”, Gustavo Alonso, Fabio Casati, Harumi Kuno, Vijay Machiraju; Springer-Verlag, Sept. 2003)
- Industry course at ETHZ, Sept. 22/23 2003 with over 40 participants from different industrial sectors.

9 Conclusion

The project is progressing according schedule. Changes in technology and specifications/standards are being accommodated during the project evolution. Problems and difficulties found during the development of the project has been identified and settled down in a timely fashion.

10 Appendices

10.1 Updated Consortiums Relevant Persons Working and/or Associated to the Project

After the contract amendment the industrial partner, HP, was substituted by Arjuna Technologies Ltd.

- UPM. The permanent staff participating in the project are: Prof. R. Jiménez-Peris, Prof. M. Patiño-Martínez.

  The currently hired staff includes: Mr. David Cadena, Mr. Alberto Erice, and Mr. David Jimenez.

  Additionally, there is a PhD student, Mr. Francisco Perez, working in the project although currently without a contractual link.

- UniBo. The permanent staff participating in the project is: Prof. Ozalp Babaoglu, Dr. Alberto Montresor, Dr. Davide Rossi.

  The hired staff is: Dr. Vance Maverick, Dr. Jaksa Vuckovic.

- ETHZ. The permanent staff participating in the project is: Prof. Gustavo Alonso.

  The currently hired staff for the project is: Mr. Daniel Joensson and Mr. Bioern Bioernstad

- McGill. The permanent staff participating in the project is: Prof. Bettina Kemme.

  The currently hired staff for the project is: Mr. Huaigu Wu and Mr. Mabrouk Chouk.

- Trieste. The permanent staff participating in the project is: Prof. Alberto Bartoli

  The hired staff for the project is: Mr. Milan Prica, Mr. Etienne Antoniutti di Muro, Mr. Cosimo Calabrese.

  Mr. Giovanni Masarin is working on the project without a contractual link, as a student.

- UNEW. Permanent staff: Prof. Santosh Shrivastava, Dr. Graham Morgan.

  Hired Staff: Mr. Simon Woodman

- Arjuna. The staff working in the project are: Dr. Stuart Wheater, Dr. Dave Ingham, Dr. Mark Little, Dr. Jim Webber.