

# XML-enabled workflow management for e-services across heterogeneous platforms

German Shegalov, Michael Gillmann, Gerhard Weikum

University of the Saarland, Department of Computer Science, P.O.Box 151150, 66041 Saarbruecken, Germany;  
E-mail: {shegalov,gillmann,weikum}@cs.uni-sb.de; http://www-dbs.cs.uni-sb.de

Edited by F. Casati, M.-C. Shan, D. Georgakopoulos. Received: 30 October 2000 / Accepted: 19 December 2000

Published online: 27 April 2001 – © Springer-Verlag

**Abstract.** Advanced e-services require efficient, flexible, and easy-to-use workflow technology that integrates well with mainstream Internet technologies such as XML and Web servers. This paper discusses an XML-enabled architecture for distributed workflow management that is implemented in the latest version of our Mentor-lite prototype system. The key asset of this architecture is an XML mediator that handles the exchange of business and flow control data between workflow and business-object servers on the one hand and client activities on the other via XML messages over http. Our implementation of the mediator has made use of Oracle's XSQL servlet. The major benefit of the advocated architecture is that it provides seamless integration of client applications into e-service workflows with scalable efficiency and very little explicit coding, in contrast to an earlier, Java-based, version of our Mentor-lite prototype that required much more code and exhibited potential performance problems.

**Key words:** Workflow management – XML/XSL – Business processes – Internet e-services – Information system interoperability

---

## 1 Introduction

### 1.1 Motivation

Advanced e-services, such as electronic auctions (e.g., Ebay), all-in-one travel planning (e.g., Expedia), automation of real-estate purchase (e.g., Realtor), or computerized court trials with electronic lawyers (still fictitious today), pose both old and new problems to the underlying software infrastructure. The new ones include, for example, the setup of behavioral contracts between services for the composition of value-added, higher-level services (see, for example, [AFH+99, BCL+00, CIJ+00, Fra99, HLG+00, LBS+99]) or the use of such services from mobile agents (see, for example, [CZB+99, KSD99, Pap99, Vei99]). Among the old, but still mostly unsolved, problems is the integration or mediation of the business objects, typically running on heterogeneous platforms, as the basic building blocks out of which e-services are composed.

The challenge in this lies in making the integration task as simple as possible, so that new e-services can be configured, deployed, and operated at very little cost (i.e., without an expensive data and system administration staff).

A business object (e.g., a purchase order or a mortgage) consists of a collection of data and some limited-scope piece of program logic (ideally both encapsulated and hidden behind an ADT interface). The classical approach for integration is to develop a unified access interface to the heterogeneous data itself through schema reconciliation and explicit data transformations (see, for example, [EP90]). However, it can be taken for granted that this kind of relatively tight integration requires intensive human efforts, and quite a few (overly) ambitious projects towards enterprise-wide data models are known as expensive failures. A much cheaper and more viable approach for integrating business objects is to wrap objects with explicit business-logic interfaces and "simply" orchestrate the control and data flow between objects at the level of an object-flow or work-flow mediator. Object-oriented middleware like CORBA, DCOM, or EJB has aimed at this kind of integration, but is more focused on wrapping the components and routing primitive requests rather than the actual mediation. The technology that really has all the ingredients to solve this integration issue is workflow management (see, for example, [GSC+99, Ley95]). A workflow consists of a set of activities, automated or intellectual/interactive ones, with explicitly specified and system-enforceable control and data flow. Each activity can in turn be viewed as invoking a method on a business object.

State-of-the-art workflow management systems, in conjunction with the underlying middleware such as CORBA, are indeed capable of integrating business objects for setting up a new e-service in amazingly short time and thus with impressively little cost (e.g., virtually no code-writing or all code automatically generated from easily constructed high-level specifications). However, the pace of the ongoing trends towards functionally richer e-services on the Internet has been so fast that the vendors of workflow management systems have had hardly any time to prepare their products for these new settings. Thus, probably no workflow product is ready for deployment on Internet-wide, arbitrarily heterogeneous platforms on a short-notice basis. The issues of how to leverage

the latest developments on XML-centered technologies and how to combine the best of several worlds into an easy-to-use and mostly administration-free infrastructure are wide open.

### 1.2 Contribution

This paper proposes an architecture that leverages XML technology for Internet-wide workflow management within advanced e-services. The key concept is to “marshal” all activity calls into special kinds of XML documents and to handle the flow between the workflow engine and the activities by some kind of XML mediator using http as the primitive transport protocol. Thus, workflow engines as well as activities and the corresponding business objects can reside on their native platforms without any special measures, and the XML mediator enables the cross-talk among all these components with all data exchanges in the form of XML documents. None of the involved parties, clients, business-object servers, or workflow servers, needs to know any details about the other parties. Clients need no software other than an Internet browser, yet handle interactions with workflows and business objects much more efficiently than with a full-blown Java implementation.

The paper presents this architecture and contrasts it with a more traditional setup that builds more heavily on CORBA (or equivalently EJB) and Java. We believe that this XML mediator approach will be beneficial as a general infrastructure and is not necessarily tied to our current context of workflow management. Its salient properties are the following:

- It builds on a rather lean software infrastructure that is virtually ubiquitous anyway: an Internet browser with simple XML/XSL support on the client side, an http server with efficient support for servlets at the mediator level, and standard SQL via JDBC and the most basic object-oriented middleware services (either CORBA or EJB) to interoperate with the various business-object and workflow servers.
- As a consequence, the entire infrastructure is very easy to set up and administer. In particular, there are no extra software installations on clients and no special requirements on the mediator’s http server; so much of the usual headache with software distribution, compatibility of software releases, and general maintenance is eliminated (or reduced to what is necessary for a ubiquitous standard infrastructure in any case).
- The mediator as a “middle man” bundles the traffic between the workflow and business-object servers on one side and (automatic as well as interactive) applications on the other side. Thus, the mediator can reuse permanently established threads and connections with the backend servers, and generally enjoys all the scalability benefits that have made three-tier architectures so prevalent on the Web. Furthermore, the computing resources of the mediator can be exploited to alleviate the load on the client or business-object server side, in adaptation to where the bottlenecks are. In particular, the mediator can push necessary transformations on the XML data that is routed between business objects and clients either to the client, making more use of XSL and browser-embedded scripts, or to the business-object servers, making more use of advanced SQL or native XML handling capabilities, or it can choose to do the bulk of this work by itself.

The proposed architecture has been fully prototyped in the latest version of our Mentor-lite distributed workflow management system [GWS+00, WGR+00]. The XML mediator itself is implemented using Oracle’s XSQL servlet, run under the Apache Web application server, with a small number of very compact so-called custom XSQL action handlers. All client functionality is embedded in an Internet browser (specifically IE5) and implemented with a fairly small amount of XSL and DHTML code.

### 1.3 Outline

The rest of the paper is organized as follows. Section 2 introduces related work. Section 3 discusses system architectures for Internet-based workflow, including the XML-enabled architecture advocated here. Section 4 discusses more specifically how the control and data flow among workflow activities is handled by the XML mediator. Section 5 presents a simple e-commerce scenario that we have implemented on our prototype system. Section 6 discusses why the proposed architecture is scalable and versatile.

## 2 Related work

Architectural issues of workflow management systems have been intensively explored in research and development (see, for example, [DKO+98, GHS95, JB96, LR99]). Much emphasis has been put on scalability and robustness to ensure industrial-strength service. Today’s most advanced products may indeed claim that they are ready for mission-critical, enterprise-wide use in terms of performance and availability. However, much of this virtue comes at the cost of a fairly large system footprint (e.g., memory requirements of the workflow system) and careful administration.

Combining workflow technology with the Internet has been successfully addressed with regard to browser-based user interfaces (e.g., for worklists) and transactional as well as “transport level” protocols (see, for example, [Moh99, SDD+97]), but there has been very little work on a deeper integration of Web and workflow technologies and their software infrastructures. In fact, the latent workflow functionality in most of today’s e-commerce applications has mostly been implemented in an ad hoc manner; typically the workflow state and context is maintained by a collection of small servlets (e.g., Active Server Pages in IIS or PHP scripts under Apache).

XML technology has been detected by the workflow community only recently. At this stage, the work in this direction appears to be limited to casting workflow specifications and system interface descriptions into XML format (see, for example, a draft of the Workflow Management Coalition [WfMC00]). This work is on the right track, but there is hardly any implementation work along these lines (one exception that we are aware of being the recent work of [CHD+99]). XML as a container for remote method invocation, with http as an underlying “transport” protocol, is intensively pursued in Microsoft’s SOAP protocol (Simple Object Access Protocol) and the corresponding efforts towards a W3C standard [SOAP]. However, SOAP aims to be a lowest common denominator among all classes of Internet applications, whereas our work focuses on richer, workflow-style, advanced e-services.

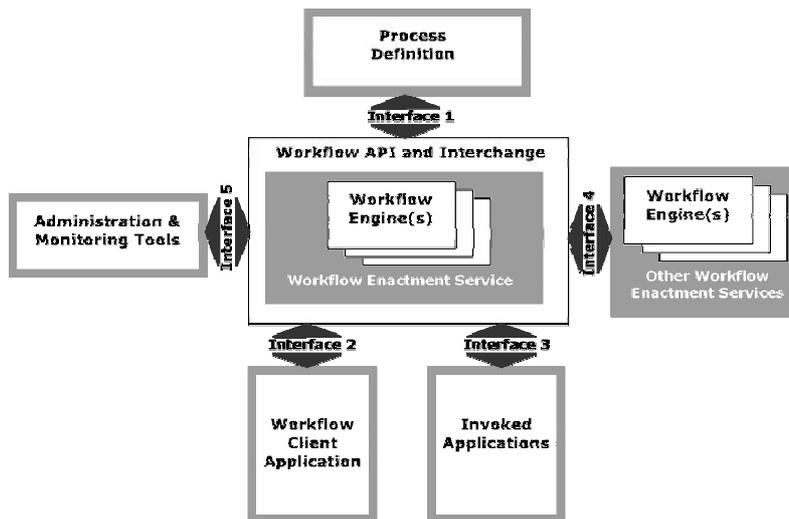


Fig. 1. Reference architecture of the WfMC

### 3 System architecture

In this section we discuss different system architectures for Internet-based workflow management. We begin, in Sect. 3.1, with the reference architecture of the Workflow Management Coalition [WfMC], an industry consortium that aims to standardize workflow-system interfaces for interoperability; this serves as a baseline against which we can compare the architectures of real systems. In Sect. 3.2 we present the original architecture of our prototype system, called Mentor-lite (i.e., an earlier version [MWG+99]). As this version of our prototype made extensive use of Java and CORBA services, we refer to it as the Java-based architecture. Only recently have we re-architected the Mentor-lite system, and the latest version is centered around XML. This is the architecture that we believe is most suitable for advanced e-services; it will be introduced in Sect. 3.3.

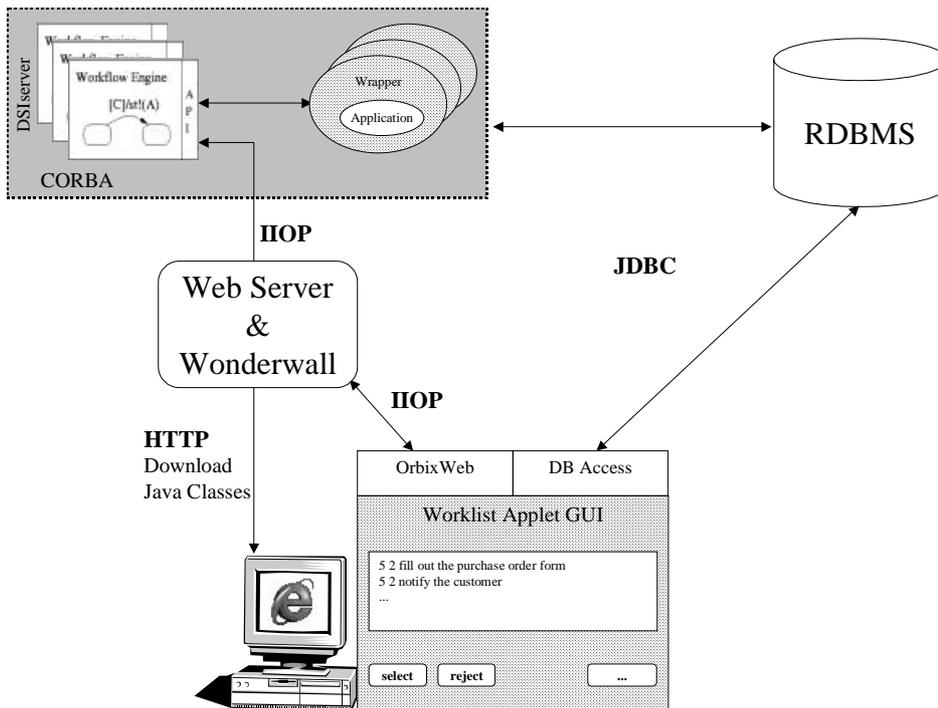
#### 3.1 Reference architecture

Workflow specification languages range from Petri-net-like or statechart-style high-level visual languages or specific types of (modal) logic all the way to scripting or simply (unstructured) collections of (ECA) rules. A workflow specification is often derived from a (business) *process modeling and definition tool* such as Aris Toolset. Upon initiation by a user a workflow specification is instantiated and interpreted by a *workflow engine*. One or more engines (e.g., on an SMP computer) form the *workflow enactment service*. During its execution a workflow spawns *activities*. These activities correspond to either *client applications* or other *invoked applications*; typically the first correspond to interactive activities such as intellectual decision making (possibly using tools such as spreadsheets) and the second to automated activities such as host applications. Workflows can spawn entire subworkflows that may be controlled by other, “external” workflow engines (running software from different vendors on servers that belong to different organizations). Finally, for analyzing workflow usage patterns and providing feedback to the business process re-engineering lifecycle, *administration and monitoring tools* are needed.

All these components and their interactions are cast into a nice framework known as the reference model of the Workflow Management Coalition [WfMC], depicted in Fig. 1. Most importantly, the WfMC has also issued standards for the five relevant APIs in this architecture. In [WfMC98] they discuss the impact of modern Internet technologies on the various APIs of the reference architecture. Obviously and most importantly, API2 should consider the fact that Internet browsers are powerful tools and already provide capabilities to embed a rich suite of client applications through plug-ins, ActiveX, Java applets, etc. As for API3, the WfMC brings up for consideration the CORBA-oriented IIOP (Internet Inter-ORB Protocol) and the services of EJB (Enterprise Java Beans) for invoking appropriately wrapped host applications. Finally, the WfMC also discusses the potential benefit of casting all workflow APIs into Java for Internet-enabling. Although these considerations by the WfMC are very useful in sorting out the various issues and establishing a common framework, they stay at an abstract level and are of very limited help in making design decisions for a concrete implementation. Most recently, the WfMC has also issued a document on Wf-XML [WfMC00] that describes how calls and parameters of the various APIs can be cast into XML format. However, how these XML messages should actually be embedded into the architecture of a real system is left wide open.

#### 3.2 Java-based architecture

Following the proliferation of Java, our first implementation of the Mentor-lite prototype [MWG+99] has made intensive use of Java on the client side and CORBA as general middleware. Our workflow engine, an interpreter for a specifically designed statechart dialect [WW97], is run as a CORBA application server, called DSIServer (for Distributed Statechart Interpreter), or as a collection of such servers on different computers with support for decentralized execution of subworkflows [GWW+00, MWW+98a, MWW+98b]. More specifically, we have used Iona’s Orbix product [IONA], one of the few industrial-strength and CORBA-compliant request brokers. The DSIServer itself makes use of an Oracle8i database for reliably tracking the workflow state and context, and also



**Fig. 2.** Architecture of the Java-based Mentor-lite version

for managing organizational and history data that is relevant for worklist policies (i.e., role resolution). This database also holds the current states of worklists and their work items.

CORBA technology is used for the API3 to invoke host applications (i.e., these applications must be wrapped with an IDL interface), and also for the API4 to other workflow engines. The latter is exploited already within Mentor-lite when workflows span multiple, organizationally separated sites each of which autonomously runs a DSIServer for the subworkflows that it is responsible for. In this case, the cross-engine communication takes place via transactional, reliable message queues that we have ourselves implemented as a specific queue server using Orbix OTS (Object Transaction Service) as a 2PC coordinator. This setup can also handle data exchanges among heterogeneous workflow servers, running engines from different vendors. However, our implementation used home-grown message formats (e.g., for signaling changes of the workflow state and context to another engine). Casting these messages into XML would allow us to coordinate workflows across heterogeneous platforms.

As for API2 between the workflow engine and the client applications and also API5 with regard to administration tools, we had chosen to make intensive use of Java applets. Worklists and their work items are presented to the user by applets embedded in the user's Internet browser. These applets are dynamically loaded from a trusted Web server. The initial applet is launched and contacts the DSIServer for work-to-do, which is implemented as an http call and handled via the IIOP-based OrbixWeb service (which needs to be installed on the client side). Once running, all applets directly access the Oracle8i database with the relevant worklist data using JDBC for performance reasons, as opposed to always having an indirection through the DSIServer (i.e., the workflow engine). When the user selects a specific work item from her/his worklist, the corresponding client application is launched as another applet that

may itself invoke standard tools (e.g., a spreadsheet program via ActiveX). This overall architecture of our first Mentor-lite version is illustrated in Fig. 2.

Based on this Java-centric version of Mentor-lite, we built a mid-sized prototype application for student enrollments, exams, etc., within our university department. The architecture proved to be viable, but we also observed a number of performance-critical issues:

1. Loading the Java applets from a trusted server into the client turned out to be time-consuming even within an Intranet (i.e., over high-speed LAN technology, but involving the usual software overhead). As these applets access resources outside the browser's sandbox by making JDBC calls, remote loading from a trusted server could not be easily eliminated. In addition, we also learned that portability and "zero-admin" installation of Java applets that include JDBC, CORBA, ActiveX, etc., calls is all but self-guaranteed because of platform-specific resource restrictions of the browser sandbox.
2. The JDBC calls themselves were relatively expensive, too. The delays due to a good number of message roundtrips were sometimes user-noticeable. In particular, the setup of a JDBC session is relatively expensive in terms of message cost.
3. With every client establishing a JDBC session with the database server, scalability may be a potential problem. Note, however, that this approach was still significantly more efficient than the alternative of involving the Orbix-based DSIServer on each and every interaction with the worklist.
4. For cross-organizational workflows over heterogeneous workflow engines, clients would have to maintain even more sessions with a suite of underlying servers. With session setup being a fairly expensive part of the protocols, even powerful clients may exhibit delays that could slow

down user interactions. Furthermore, usually all TCP ports other than the http port 80 are blocked by firewalls. Thus, we inevitably have to use http for Internet-wide interoperability. With this constraint, XML is the best choice as a data exchange format. The fact that XML is a rather verbose and long message format can be easily compensated by appropriate compression techniques (see, for example, [LS00]).

### 3.3 XML-based architecture

Based on the lessons learned with the Java-based architecture, we redesigned the interfaces to client applications (API2) and invoked applications (API3) and modified the affected components of the Mentor-lite prototype system, leading to what we refer to as the XML-based version of Mentor-lite.

In essence, we modified the messages that are received and sent by clients and invoked applications so that they are in XML format, using a set of Mentor-lite-specific XML tags. Messages now contain the activity name and workflow id as an XML element, and further XML elements for input and output parameters. For readability and easier processing we discriminated parameters into two classes:

- *Business data* are input or output parameters that influence the persistent business objects on which the activity's application operates. For example, when invoking an activity to receive a new order, the customer id, id of the ordered product, and the ordered quantity are business data.
- *Flow control data* are input or output parameters that "merely" serve to drive the control (and data) flow among activities in the workflow, but have no long-term impact on persistent business objects beyond the scope of the current workflow. This category especially includes return codes of activities.

We re-implemented the client software so that it makes use of the browser's capabilities for local XML processing (using IE5 in our implementation). The client can filter and sort data through XSL stylesheets, which is exploited on worklists, and DHTML (Dynamic HTML) is used for presenting XML data in a user-friendly GUI. This modification led to a drastic reduction of the client code and eliminated performance problem 1 of the earlier Java-based architecture (see Sect. 3.2). The client code for the student enrollment and exams application mentioned in Sect. 3.2 was reduced to a few hundred lines of JavaScript/JScript and XSL stylesheets.

To address the other performance problems mentioned in Sect. 3.2 we introduced an *XML mediator* as a "middle man" between the clients and the workflow engines as well as the Oracle8i database server. This mediator was designed as a servlet that can run in any http server. Its purpose is to feed XML data to the client activities and invoked applications and collect their XML output, while being able to communicate with the backend servers much more efficiently. In particular, the mediator servlet can maintain a moderate number of permanently open JDBC sessions with the Oracle8i database server, which eliminates performance problems 2 and 3 to a large extent (see Sect. 3.2).

As for the mediator implementation we massively benefited from the recent release of Oracle's XSQL servlet [XSQL]

(which was still in Beta status at the time when we carried out our implementation). This service, which is implemented in Java and can run in any standard http/servlet server (e.g., Apache, which is used in our prototype), mediates between the SQL and the XML worlds by converting the results of SQL statements into XML and generating SQL calls to store incoming XML data via JDBC. Most importantly, XSQL could be easily extended through so-called custom XSQL action handlers to accommodate functions that are specific to our workflow environment. We will present details of this customized extension in Sects. 4 and 5.

The XML mediator is also the ideal middle man in a cross-organizational workflow with highly heterogeneous server platforms, as it encapsulates the servers with regard to the clients and can maintain a moderate number of sessions with heterogeneous workflow or business-object servers with reasonable efficiency. Thus, problem 4 mentioned in Sect. 3.2 is also, to a large extent, rectified in the new architecture.

The overall XML-based architecture of Mentor-lite is illustrated in Fig. 3.

## 4 The XML mediator

### 4.1 Conceptual overview

The XML mediator provides a high-level abstraction for WFMS interoperability in the context of global (i.e., cross-enterprise) workflow management. It offers an Internet-based, worldwide accessible interface to invoke workflows or activities and pass results to the proper destinations. By viewing an activity as an abstraction of either an invoked application or a subworkflow the mediator enables seamless integration of all kinds of business activities. As underlying IT infrastructure we consider workflow, application, and business-object servers, where the latter are typically database servers. Our goal is to hide all infrastructure details (location of servers, etc.) behind a simple XML interface.

Consider a scenario, depicted in Fig. 4, with two enterprises involved in a global business process, i.e., an e-service such as providing personalized multimedia newspapers to subscribers. Both business process participants, BPP1, which could be the business portal for this e-service, and BPP2, which could be the content provider, have their own specific IT solutions. The XML mediator facilitates the interoperation between these two information systems, thus making global workflows feasible (e.g., for setting up the service of a new subscriber). There may also be interactive activities, which are introduced solely for the purpose of the global workflow. The XML mediator hides implementation details (underlying middleware, firewalls, etc.) of BPP1 and BPP2 to each party and to the interactive activity. Special XML messages that contain only business and flow control data are exchanged between the various parties via the XML mediator. The mediator is responsible for the message delivery according to the receiver's identification and the tags in the XML message. In our prototype, for example, different tags are used to identify business versus flow control data, with the latter being delivered to the workflow engine and the business data sent to an activity or, upon completion of the activity, back to a database server. Note that an activity is itself an abstraction that comprises

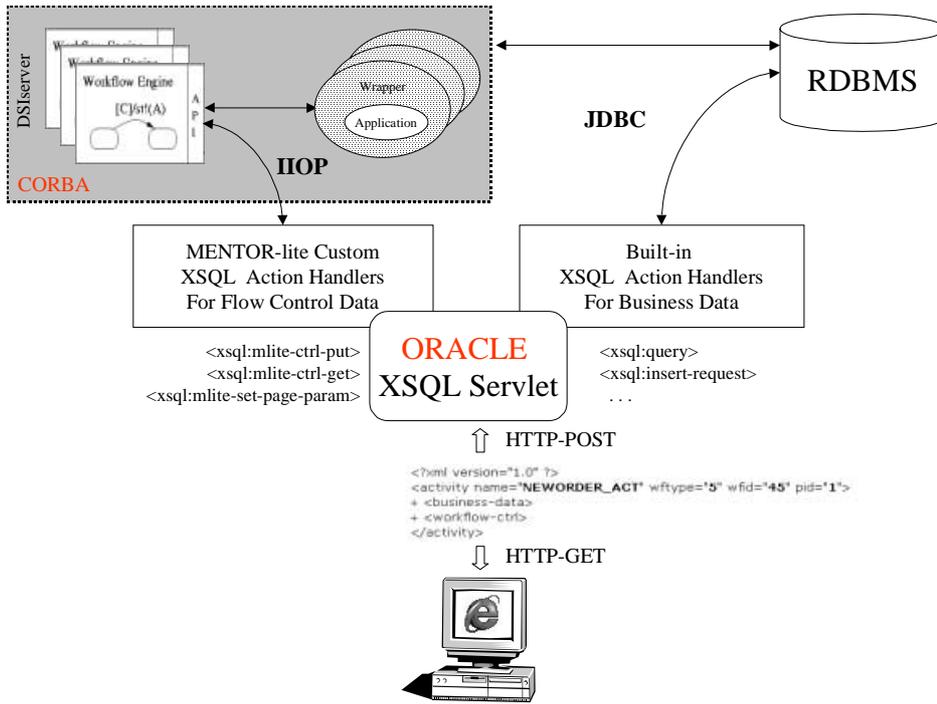


Fig. 3. Architecture of the XML-based Mentor-lite version

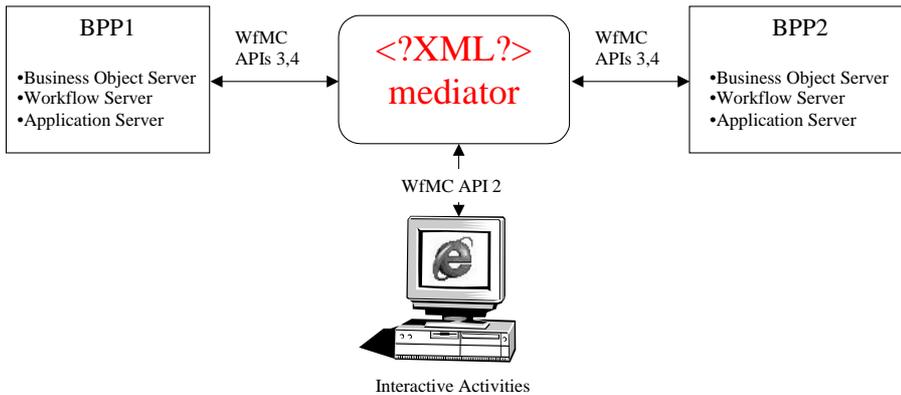


Fig. 4. Role of the XML mediator

interactive activities (the case on which we mostly focus), automated activities that invoke applications, or activities that encapsulate a subworkflow running on a different workflow engine.

We identify activities, which may be subworkflows, by a triple  $\langle \text{workflow-type, workflow-id, process-id} \rangle$ . XML messages that are sent to an activity to invoke the activity (and whatever application or subworkflow is behind the activity) are structured as follows:

```
<?xml version='1.0' ?>
<activity name='some name' wftype='x' wfid='y' pid='z'>
  <business-data>
    <relationl><rowl><attr1>...</attr2></rowl> ...
  </relationl>
  ...
</business-data>
  <workflow-ctrl>
    <variable name='var-name1' value='var-value1' />
    ...
  </workflow-ctrl>
</activity>
```

Here we assume that business objects typically reside in relational databases, but extensions to incorporate object-oriented or object-relational databases would be straightforward.

```
1 public interface XSQLActionHandler {
2   void init(XSQLPageRequest env, Element actionElement) ... ;
3   void handleAction(Node rootNode) ... ;
4 }
```

Fig. 5. Generic Java interface of the XSQL action handlers

When an activity completes, it sends an analogously structured XML message with its output to the mediator, which then parses the message and derives the necessary follow-up actions such as storing new or modified business objects and sending flow control data to a workflow engine.

#### 4.2 Interfaces

Upon the start of an activity, the client's Web browser obtains its input data in XML format by an http get call. Analogously, results from the activity are returned by an http post call. The corresponding DSIserver to which the output should eventu-

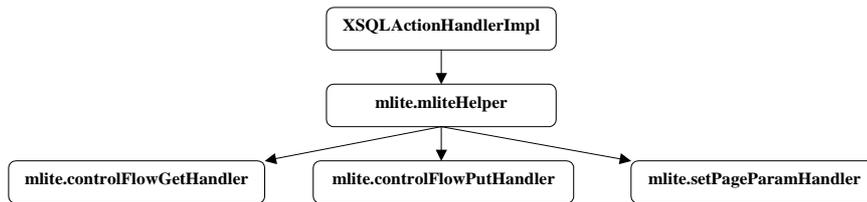


Fig. 6. Class hierarchy of the action handlers

Table 1. Mentor-lite specific custom XSQL action handlers

Action element	Action handler class	Action
<xsql:mlite-ctrl-get>	mlite.controlFlowGetHandler	to read flow data
<xsql:mlite-ctrl-put>	mlite.controlFlowPutHandler	to write flow data
<xsql:mlite-set-page-param>	mlite.setPageParamHandler	to read flow data to be set as a parameter at the page-level

ally be delivered is referenced by the triple <workflow type, workflow ID, process ID> which is automatically resolved by the XML mediator using the CORBA name service. Therefore, the application that implements the activity on the client side need not know details about where input data comes from and how result data is handled further on.

The core of the XML mediator is implemented using Oracle's XSQL servlet. This Java servlet serves to combine results of SQL queries into XML documents or to extract SQL updates from such documents. A strong feature of the XSQL servlet is the option for adding "custom XSQL action handlers" that can be invoked by the XSQL servlet based on special XML elements, marked by the prefix *xsql*, in the XSQL page. Both custom and built-in action handlers conform to the Java interface shown in Fig. 5. The function *init* (line 2) serves to initialize the action handler, typically by reading the attributes of the action element. The *handleAction* function (line 3) executes the action itself. The parameter *rootNode* refers to the node that will be inserted in place of the action element in the resulting XML document. For each appearance of an XSQL action element in the XSQL page, the XSQL servlet creates a new instance of the corresponding action handler class, and invokes the method *init* and *handleAction* on the recently created instance, subsequently [XSQL].

Custom action handlers are needed in the context of workflow management to handle the flow control and business data that are exchanged between an activity and the workflow engine. Table 1 shows three basic XML elements and corresponding action handlers for reading and manipulating control flow variables, i.e., variables that are required by the workflow engine for control flow decisions (e.g., return codes).

To keep the implementation of the custom action handlers simple, we have designed a class hierarchy shown in Fig. 6. All action handler classes are derived from the basic class *mliteHelper* that includes generic functions like communication interfaces (e.g., requests to the CORBA name service). The class *mliteHelper* itself extends the *XSQLActionHandlerImpl* class, a base implementation of XSQL action handlers that includes a set of useful helper methods (e.g., for error handling and monitoring). The *XSQLActionHandlerImpl* class is part of the runtime library by Oracle.

### 4.3 Implementation

Figure 7 shows the implementation of the *mliteHelper* class that extends the base implementation of the *XSQLActionHandlerImpl* class. It provides a simple API to identify the workflow engine responsible for the current workflow instance and the primary error handling that is necessary for using the CORBA name service.

The lines 4–17 initialize the CORBA interface and connect to the name service. The function *getWfEngine* (lines 18–25) determines the workflow engine that is responsible for the execution of the current workflow instance. Thus, the application programs that implement the workflow activities do not need to know anything about the location of the engine or its distribution over the network. The function *insertException* (lines 26–29) supports exception handling.

As an example of the implementation of a workflow-specific action handler, Fig. 8 shows the Java code for the *controlFlowPutHandler* class that is associated with the action element <xsql:mlite-ctrl-put> and will be automatically called by the XSQL servlet when processing the XML output of an activity. The *controlFlowPutHandler* class overwrites the method *handleAction* of the base implementation. When *handleAction* is invoked, we obtain a reference to the posted XML document as an XML DOM object by calling the method *getPostedDocument* of the interface *XSQLPageRequest* in line 3. Subsequently, we are able to manipulate or read nodes of the posted XML document in lines 4–7. First we extract the attributes *wftype*, *wfid*, and *pid* from the root node *activity*, i.e., the responsible DSIserver can be identified by calling the method *getWfEngine* of the base class *mliteHelper*. The list of the control flow variables is obtained by calling the *getElementsByTagName* (lines 8–11) on the workflow-ctrl DOM element object. After the connection to the workflow engine is established (line 18), flow control data can be transferred to the workflow engine by calling *wfengine.put* method (lines 21–28).

## 5 A simple case study

In this section, we present the implementation of a simplified e-commerce scenario as an example of a workflow application. This case study serves as a proof of concept for our mediator approach. In addition, it shows that the implementation of

```

1 public abstract class mliteHelper extends XSQLActionHandlerImpl {
2     private static NamingContext rootContext;
3     private static org.omg.CORBA.ORB orb;
4     static {
5         System.err.print("initializing CORBA interface ... ");
6         orb = ORB.init((String[]) null, null);
7         System.err.println("done");
8         System.err.print("initializing NamingService ... ");
9         try {
10            org.omg.CORBA.Object ns =
11                orb.resolve_initial_references("NameService");
12            rootContext = NamingContextHelper.narrow(ns);
13        } catch(Exception e) {
14            System.err.println(e.getMessage());
15        }
16        System.err.println("done");
17    }
18    public dsiserver getWfEngine(String engineName) throws Exception {
19        System.err.print("Search for the engine " + engineName + " ... ");
20        NameComponent[] name = new NameComponent[1];
21        name[0] = new NameComponent(engineName, null);
22        org.omg.CORBA.Object wfe = rootContext.resolve(name);
23        System.err.println("done");
24        return mlite.dsiserverHelper.narrow(wfe);
25    }
26    public void insertException(org.w3c.dom.Node n, Exception e) {
27        this.reportError(n, e.toString());
28        e.printStackTrace();
29    }
30 }

```

**Fig. 7.** Implementation of action handler *mliteHelper*

```

/* error handling code isn't shown here */
1 public class controlFlowPutHandler extends mliteHelper {
2     public void handleAction( Node rootNode ) throws SQLException {
3         Document xmlDoc = this.getPageRequest().getPostedDocument();
4         Element xmlRoot = xmlDoc.getDocumentElement();
5         String wftype = xmlRoot.getAttribute("wftype");
6         String wfid = xmlRoot.getAttribute("wfid");
7         String pid = xmlRoot.getAttribute("pid");
8         NodeList wfCtrlBlock =
9             xmlRoot.getElementsByTagName("workflow-ctrl");
10        NodeList variables =
11            ((Element) wfCtrlBlock.item(0)).getElementsByTagName("variable");
12        int varcount = variables.getLength();
13        String engine_name = "dsiserver" +
14            wftype + '_' +
15            wfid + '_' +
16            pid;
17        try {
18            dsiserver wfengine = this.getWfEngine(engine_name);
19            Element curr_elem;
20            String varval;
21            for(int i=0; i < varcount; i++) {
22                curr_elem = (Element) variables.item(i);
23                varval =
24                    curr_elem.getAttribute("name") +
25                    '=' +
26                    curr_elem.getAttribute("value");
27                wfengine.put(varval, "");
28            }
29            wfengine._release();
30            this.reportStatus(rootNode,
31                "result",
32                varcount + " variables inserted");
33        } catch(Exception e) {
34            insertException(rootNode,e);
35        };
36    }
37 }

```

**Fig. 8.** Implementation of action handler *controlFlowPutHandler*

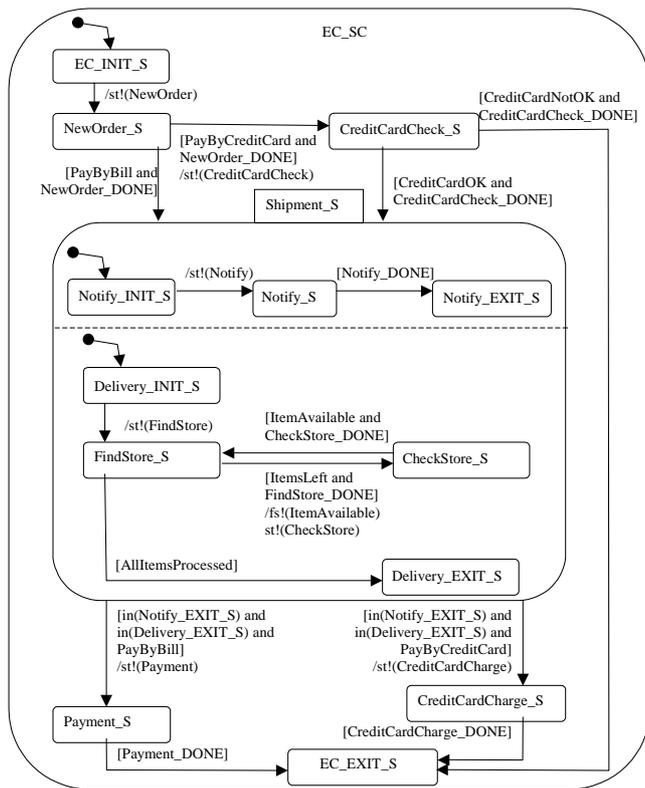


Fig. 9. State chart of the *electronic commerce (EC)* workflow example

new workflow applications for e-services can be carried out in a straightforward manner in very short time and thus with very low cost. In the example, we focus on the interface between the workflow engines and the applications invoked for interactive activities.

The workflow is based on the TPC-C order-entry benchmark for transaction systems [TPC], with the key difference that we combine multiple transaction types into a workflow and further enhance the functionality (see [GMW+99] for a full description of this workflow).

Figures 9 and 10 show the workflow specification as a state and activity chart, a formalism [Har87, HG97] that has been adopted for the behavioral dimension of the UML industry standard and is used in our prototype system Mentorlite [WW97, MWW+98a]. Each state in Fig. 9 corresponds to an activity in Fig. 10 or one (or multiple, parallel) subworkflow(s), except for initial and final states. We assume that for every activity *act* the condition *act\_DONE* is set to true when *act* is finished.

The workflow proceeds as follows: initially, the *NewOrder* activity is started. After the termination of *NewOrder*, the control flow branches. If the customer wants to pay by credit card, the condition *PayByCreditCard* is set and the *CreditCardCheck* activity checks the validity of the credit card. If there are problems with the credit card, the workflow is terminated. Otherwise the shipment, represented by the nested top-level state *Shipment.S*, is initiated spawning two orthogonal/parallel subworkflows. The first subworkflow has only one activity that sends a notification mail. The second subworkflow (sequentially) invokes for each ordered item an activity that identifies a store from which the item could be shipped. Then,

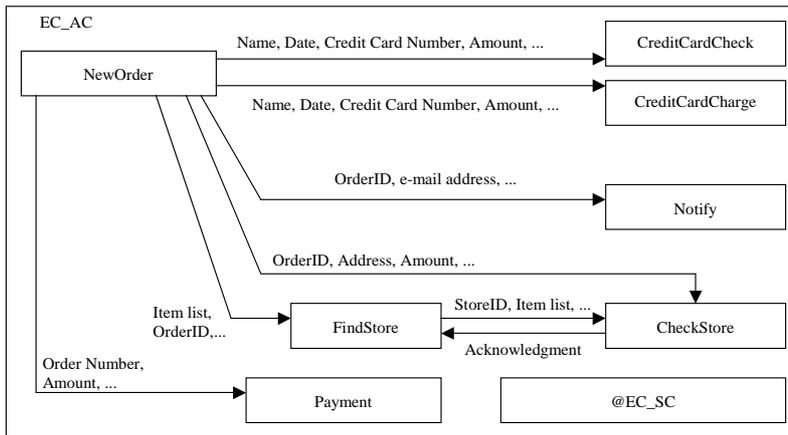
a second activity instructs the store to deliver the item and waits for an acknowledgement. The two activities *FindStore* and *CheckStore* are repeated within a loop over all ordered items. After the termination of both subworkflows, the control flow is synchronized, and branches again depending on the mode of payment. The workflow terminates in the state *EC\_EXIT.S*.

In the following, we look closer at the *NewOrder* activity to illustrate the role of the XML mediator in the activity handling. This activity is implemented as a Dynamic HTML (DHTML) application (see Fig. 11). In general, there are two approaches to handle XML on the client side using Microsoft IE 5.0. First, the XML data can be transformed by the stylesheet engine of the browser into DHTML after downloading XML and XSL. Alternatively, a client application can be implemented in DHTML directly. In the latter case, so-called “XML data islands” [XML] can be defined within a DHTML document, in order to separate the input data from the presentation. An XML data island can easily be bound (“Data Binding” [XML]) to an HTML element as a dynamic source object (DSO) to be visualized within the DHTML document. The main advantage of using this Data Binding is that the HTML element will be updated automatically if the corresponding DSO has changed. Once initialized in the browser, the DHTML document for the *NewOrder* activity contains an empty XML data island `<xml id=“dsoCustomer”></xml>`. The `<xml id ... >` tag is a special HTML tag for defining XML data islands. For convenience the user may simply input the customer id and click on the “Check This” button. On this event the function *checkCustomer()* (Fig. 12) will be executed, which loads the content of the XSQL page *customer.xsql* (Fig. 13), processed by XSQL servlet, into the XML DOM object *dsoCustomer* (line 4). Form fields like “Phone” are automatically filled with the content of the `<phone>` tag of *dsoCustomer* (line 6 in Fig. 14) by specifying `<input size=“50” name=“phn” datasrc=“# dsoCustomer” datafld=“phone”>` in the DHTML script. The attributes *datasrc* and *datafld* refer to the DSO object and the field of interest in this object. All other form fields are filled analogously (lines 3–14).

Unlike customer data, items can be added to the “shopping cart” several times. To avoid repeatedly checking item IDs by calling the XML mediator, we prefetch relevant parts of the product catalog into the browser by defining an XML data island. All manipulations on the DHTML page, such as adding items, are processed locally without involving any servers or the XML mediator.

When the *NewOrder* activity is finished, i.e., the user has entered all ordered items and pushed the submit button, an XML document that includes both the business and flow control data is sent to the XSQL servlet by a http post call. An example for such an output is given in Fig. 15. The structure of the output is generic and the same for all activities. The `<activity>` element in line 2 specifies the activity and the corresponding workflow instance followed by a block of business data (lines 3–18) and a block of flow control data (lines 19–23).

The XSQL page that receives the document is shown in Fig. 16. The `<xsql:insert-request>` action element (lines 3–5), one of the built-in action elements from Oracle, refers to an XSQL action handler that stores the business data in the specified table. The action element `<xsql:mlite-ctrl-put>` (line 6) is assigned to the action handler *controlFlowPutHandler* that



**Fig. 10.** Activity chart of the *electronic commerce (EC)* workflow example

**Fig. 11.** Output of the DHTML application for activity *NewOrder*

```

1 function checkCustomer() {
2   dsoCustomer.async = false;
3   url = "customer.xsql?customerID=" + customer.ID.value;
4   dsoCustomer.load(url);
5 }

```

**Fig. 12.** JScript function for fetching the customer data according to the customer ID (part of the DHTML code for activity *NewOrder*)

```

1 <?xml version="1.0" ?>
2 <customer connection="mlite-demo">
3   <xsql:query xmlns:xsql="urn:oracle-xsql"
4     rowset-element="" id-attribute=""
5     row-element="" tag-case="lower">
6     SELECT *
7     FROM CUSTOMERS
8     WHERE id={ @customerID }
9   </xsql:query>
10 </customer>

```

**Fig. 13.** XSQL page customer.xsql

transfers the flow control data to the DSIServer, i.e., the workflow engine, as described in Sect. 4.

The bottom line of this walk-through of our example scenario is that surprisingly little code is needed to build a simple but not unrealistic e-service application.

```

1 <?xml version="1.0" ?>
2 <customer>
3   <id>2</id>
4   <name>Abraham O'Brian</name>
5   <mail>a.obrian@mailbox.com</mail>
6   <phone>+1(802)1234567890</phone>
7   <address>
8     <street>3384 Alohea Ave</street>
9     <city>Honolulu</city>
10    <state>HI</state>
11    <zip>96816-2202</zip>
12  </address>
13  <balance>0</balance>
14  <discount>0</discount>
15 </customer>

```

**Fig. 14.** Output produced by XSQL servlet after processing *customer.xsql*

## 6 Discussion of scalability and versatility

As pointed out earlier, the presented XML-enabled, mediator-based architecture is both scalable and highly versatile. Here, scalability means the ability to grow the entire system by adding clients, workflow servers, or business object servers without creating any performance bottlenecks. By versatility, we refer to the support of a large diversity of components including heterogeneous workflow engines and application servers.

The scalability of the mediator-based architecture draws from the following two salient properties. First, the mediator removes from the clients all resource-critical functionality that deals with server connections; such functions would otherwise not scale up with an increasing number of clients. In particular, the mediator uses the standard technique of reusing server connections, for example, JDBC sessions, with workflow or business object servers from a connection pool, rather than opening a new connection upon each client request. Second, the mediator is stateless; so it does not have to remember anything about clients or servers beyond the scope of a single http get or post request. This results in very little overhead and high mediator throughput. Furthermore, and most importantly, the mediator can be easily replicated at multiple sites, similar to the Web content servers operated by application service providers. This ensures high availability of the medi-

```

1 <?xml version="1.0"?>
2 <activity name="NewOrder" wftype="5" wfid="45" pid="1">
3   <business-data>
4     <order>
5       <row>
6         <order_id>1</order_id>
7         <customer_id>3</customer_id>
8         <prod_id>100223</prod_id>
9         <count>12</count>
10      </row>
11     <row>
12       <order_id>1</order_id>
13       <customer_id>3</customer_id>
14       <prod_id>100002</prod_id>
15       <count>1</count>
16     </row>
17   </order>
18 </business-data>
19 <workflow-ctrl>
20   <variable name="PayByCreditCard" value="1"/>
21   <variable name="PayByBill" value="0"/>
22   <variable name="NEWORDER_OK" value="1"/>
23 </workflow-ctrl>
24 </activity>

```

Fig. 15. XML output of the activity *NewOrder*

```

1 <?xml version="1.0" encoding="iso-8859-1"?>
2 <activity xmlns:xsql="urn:oracle-xsql" connection="mlite-demo">
3   <xsql:insert-request
4     tag-case="lower"
5     table="orders"/>
6   <xsql:mlite-ctrl-put/>
7 </activity>

```

Fig. 16. NewOrder.xsql for http post

ator service, and standard techniques can be applied for load balancing across the various mediator instances.

We have emphasized in this paper the XML mediator's role in encapsulating servers vis-a-vis e-service clients and simplifying the clients themselves; this aspect has primarily driven our re-implementation of Mentor-lite. However, the developed architecture is much more versatile and can be leveraged in further ways as follows: when a client invokes an application, by getting or posting an appropriate XML document from/to the mediator, it does not need to know whether this is a stand-alone application that runs on some business object server or a full-fledged (sub-) workflow or even an interactive activity that will be added to the worklist of another client. All these different kinds of application types are encapsulated by the mediator. The client software merely needs to understand the application's interface and be able to interpret its input and output parameters. Thus the mediator unifies the interfaces to all types of activities that constitute an e-service. Note that this is a major simplification compared to the various, fairly sophisticated, interfaces devised by the Workflow Management Coalition. Of course, the XML mediator itself needs to cope with the full heterogeneity of different workflow engines, CORBA or EJB servers, database servers, Web servers with servlet engines, etc., but our approach concentrates this complexity in one piece of code, thus simplifying system maintenance and e-service administration.

## 7 Concluding remarks

In this paper we have described an XML-enabled architecture for distributed workflow management on the Internet. The architecture is fully implemented in our Mentor-lite prototype system. In comparison to our earlier, Java-centric, implementation, we have achieved both substantial performance improvements and also a drastic reduction in the coding that is necessary for building workflow applications. Especially for the latter reason we believe that this architecture, that builds on ubiquitous standard infrastructure enhanced by an XML mediator, is particularly well geared for easy, fast, and inexpensive setup of advanced e-services on the Internet.

To this end, it is crucial that clients do not need any specific software installations or setups, as it would be a nightmare to make millions of potential clients ready for e-service usage and maintain their local software environment as services undergo evolution. In this regard, Java applets, upon which our earlier prototype relied, are much less ubiquitous than the marketing leads us to believe, once these applets include non-trivial resource manipulation such as IIOP calls. Our current implementation, as presented in this paper, avoids these complications by delegating these non-trivial resource accesses to the XML mediator as an interface between workflow engines and activities. This way, XML technology allows clients to participate in Internet workflows and e-services without any special setup and with user-acceptable performance.

Future work includes the practical proof of concept that the mediator can also be leveraged for the communication between different WFMSs (API 4 of the WfMC reference architecture). Furthermore, we plan to extend the mediator to support transactional communication (i.e., using 2PC among workflow engines, applications, and business-object servers). Thus, in the long term, the mediator should evolve into a standard middle tier for data exchange between WFMSs and applications as well as between different WFMSs.

## References

- [AFH+99] Alonso G, Fiedler U, Hagen C, Lazcano A, Schuldt H, Weiler N (1999) WISE: business to business e-commerce. In: Proc. Int. Workshop on Research Issues in Data Engineering (RIDE), Sydney, Australia
- [BCL+00] Bornhövd C, Cilia M, Liebig C, Buchmann A (2000) An infrastructure for meta-auctions. In: Proc. Int. Workshop on Advanced issues of E-Commerce and Web-Based Information Systems (WECWIS), San Jose, California
- [CHD+99] Chen Q, Hsu M, Dayal U, Griss M (1999) Multi-agent cooperation, dynamic workflow and xml for e-commerce automation. Technical Report, Hewlett Packard Software Technology Laboratory 1999. <http://www.hpl.hp.com/techreports/1999/HPL-1999-136.html>
- [CIJ+00] Casati F, Ilnicki S, Jin L, Krishnamoorthy V, Shan M.-C (2000) eFlow: a platform for developing and managing composite e-services. Technical Report, Hewlett Packard Software Technology Laboratory 2000. <http://www.hpl.hp.com/techreports/2000/HPL-2000-36.html>
- [CZB+99] Chrysanthis P, Znati T, Bunjeree S, Chang S.-K (1999) Establishing virtual enterprises by means of mobile agents. In: Proc. Int. Workshop on Research Issues in Data Engineering (RIDE), Sydney, Australia
- [DKO+98] Dogac A, Kalinichenko L, Tamer Ozsu M, Sheth A (eds) (1998) Workflow management systems and interoperability. NATO Advanced Study Institute, Springer, Berlin Heidelberg New York
- [EP90] Elmagarmid AK, Pu C (guest editors) (1990) Special Issue on Heterogeneous Databases, ACM Comput Surv 22(3)
- [Fra99] Fraternali P (1999) Tools and approaches for developing data-intensive web applications: a survey. ACM Comput Surv 31(3)
- [GHS95] Georgakopoulos D, Hornick M, Sheth A (1995) An overview of workflow management: from process modeling to workflow automation infrastructure. Distrib Parallel Databases 3(2)
- [GMW+99] Gillmann M, Muth P, Weikum G, Weissenfels J (1999) Benchmarking of workflow management systems. In: Proc. of German Conf. on Database Systems in Office, Engineering, and Scientific Applications (BTW), Freiburg, Germany
- [GSC+99] Georgakopoulos D, Schuster H, Cichocki A, Baker D (1999) Managing process and service fusion in virtual enterprises. Inf. Syst. 24(6)
- [GWS+00] Gillmann M, Weissenfels J, Shegalov G, Wonner W, Weikum G (2000) A goal-driven auto-configuration tool for the distributed workflow management system Mentor-lite. In: Proc. of ACM SIGMOD Conf. on Management of Data (SIGMOD), Dallas, Texas
- [GWW+00] Gillmann M, Weissenfels J, Weikum G, Kraiss A (2000) Performance and availability assessment for the configuration of distributed workflow management systems. In: Proc. Int. Conf. on Extending Database Technology (EDBT), Konstanz, Germany
- [Har87] Harel D (1987) State charts: A visual formalism for complex systems. Sci Comput Program 8
- [HG97] Harel D, Gery E (1997) Executable object modeling with statecharts. IEEE Comput 30(7)
- [HLG+00] Hoffner Y, Ludwig H, Gülcü C, Grefen P (2000) Architecture for cross-organisational business processes. In: Proc. Int. Workshop on Advanced issues of E-Commerce and Web-Based Information Systems (WECWIS), San Jose, California
- [IONA] IONA Technologies. <http://www.iona.com>
- [JB96] Jablonski S, Bussler C (1996) Workflow-management, modeling concepts, architecture and implementation. International Thomson Computer, London
- [KSD99] Kaiser G, Stone A, Dossick S (1999) A mobile agent approach to lightweight process workflow. In: Proc. Int. Process Technology Workshop (IPTW), Villars de Lans, France
- [LBS+99] Ludwig H, Bussler C, Shan M.-C, Grefen P (1999) Cross-organizational workflow management and coordination. WACC'99 Workshop Report. <http://www.zurich.ibm.com/%7Ehlu/WACCworkshop/Summary.html>
- [Ley95] Leymann F (1995) Workflows make objects really useful. In: Proc. Int. Workshop on High Performance Transaction Systems (HPTS)
- [LR99] Leymann F, Roller D (1999) Production workflow: concepts and techniques. Prentice-Hall, Englewood Cliffs, N.J., USA
- [LS00] Liefke H, Suciu D (2000) Xmill: an efficient compressor for XML data. In: Proc. of ACM SIGMOD Conf. on Management of Data (SIGMOD), Dallas, Texas
- [Moh99] Mohan C (1999) Workflow management in the internet age. Tutorial at Workshop on Next Generation Information Technologies and Systems (NGITS), Zikhron-Yaakov, Israel. <http://www-rodin.inria.fr/~mohan>
- [MWG+99] Muth P, Weissenfels J, Gillmann M, Weikum G (1999) Integrating light-weight workflow management systems within existing business environments. In: Proc. Int. Conf. on Data Engineering (ICDE), Sydney, Australia
- [MWW+98a] Muth P, Wodtke D, Weissenfels J, Weikum G, Kotz Dittrich A (1998) Enterprise-wide workflow management based on state and activity charts. In: [DKO+98]
- [MWW+98b] Muth P, Wodtke D, Weissenfels J, Kotz Dittrich A, Weikum G (1998) From centralized workflow specification to distributed workflow execution. Intell Inf Syst 10(2)
- [Pap99] Papazoglou M (1999) The role of agent technology in business to business electronic commerce. In: Klusch M, Shehory O, Weiß G. (eds) Cooperative Information Agents III, Lecture Notes in Computer Science, vol. 1652. Springer, Berlin Heidelberg New York
- [SDD+97] Shan M.-C, Davis J, Du W, Huang Y (1997) HP Workflow research: past, present, and future. Technical Report, Hewlett Packard Software Technology Laboratory. <http://www.hpl.hp.com/techreports/97/HPL-97-105.html>

- [SOAP] Microsoft Developer Network (MSDN) (2000) Simple Object Access Protocol. <http://msdn.microsoft.com/workshop/xml/general/soapspec.asp>
- [TPC] Transaction Processing Performance Council. <http://www.tpc.org>
- [Vei99] Veijalainen J (1999) Transactions in mobile electronic commerce. In: Saake G, Schwarz K, Türker C (eds) Transactions and database dynamics, Lecture Notes in Computer Science, vol. 1773. Springer, Berlin Heidelberg New York
- [WfMC] Workflow Management Coalition (WfMC) (1999) Reference model and glossary. <http://www.wfmc.org>
- [WfMC98] Workflow Management Coalition (WfMC) (1998) Workflow and internet: catalysts for radical change. White Paper, June 1998. <http://www.wfmc.org>
- [WfMC00] Workflow Management Coalition (WfMC) (2000) Workflow standard - interoperability Wf-XML binding. Document Draft, January 2000. <http://www.wfmc.org>
- [WGR+00] Weissenfels J, Gillmann M, Roth O, Shegalov G, Wöner W (2000) The Mentor-lite prototype: a light-weight workflow management system. In Proc. Int. Conf. on Data Engineering (ICDE), San Diego, California
- [WW97] Wodtke D, Weikum G (1997) A formal foundation for distributed workflow execution based on state charts. In Proc. Int. Conf. on Database Theory (ICDT), Delphi, Greece
- [XML] Microsoft Developer Network (MSDN) (2000) XML Developer Center. <http://msdn.microsoft.com/xml/default.asp>
- [XSQL] Oracle Technology Network. <http://technet.oracle.com/tech/xml/xsql.servlet>