



Workflow Management: State of the Art vs. State of the Products

Weimin Du, Ahmed Elmagarmid*
Software Technology Laboratory
HPL-97-90
July, 1997

workflow

It has been over ten years since the first workflow product was introduced. Despite the large number of workflow vendors and research efforts all over the world, despite all the rosy predictions about the workflow market, and despite all the hype about the workflow technology, workflow is still far from pervasive. This paper tries to understand the situation from a technical point of view, focusing on the development and enactment aspects of workflow processes. We discuss the current capabilities of workflow products, major issues that need to be addressed before workflow can be pervasive, as well as possible future trends and research that will help workflow succeed.

Internal Accession Date Only

*Purdue University, West Lafayette, Indiana

To be published in the proceedings of *NATO Advanced Study Institute on Workflow Management Systems*, Istanbul, Turkey August 12-21, 1997.

© Copyright Hewlett-Packard Company 1997

Workflow Management: State of the Art vs. State of the Products

Weimin Du

Hewlett-Packard Labs.
Palo Alto, CA 94304, USA

Ahmed Elmagarmid

Purdue University
West Lafayette, IN 47907, USA

Abstract

It has been over ten years since the first workflow product was introduced. Despite the large number of workflow vendors and research efforts all over the world, despite all the rosy predictions about the workflow market, and despite all the hype about the workflow technology, workflow is still far from pervasive. This paper tries to understand the situation from a technical point of view, focusing on the development and enactment aspects of workflow processes. We discuss the current capabilities of workflow products, major issues that need to be addressed before workflow can be pervasive, as well as possible future trends and research that will help workflow succeed.

1 Introduction

1.1 Background

Computer technology, as a major industry revolution, has evolved to a point that it has been successfully used in all application domains such as banking, finance, and telecommunication. The application of computer technology has greatly increased productivity and has provided better services. Despite their great success, the early version of computer applications have the following two major drawbacks.

First, they are monolithic in nature. All business policies and information accesses were hard-coded into the applications. These systems are hard to maintain and enhance, when the business policies and data changed. The advance of database technology has successfully separated data accesses from the applications. As database applications, they are more adaptive to data changes. On the other hand, business policies are still hard-coded and any changes would require modifying application code.

Second, they are isolated. The computer applications (especially those developed in old days) are usually designed and developed to work independently to solve specific problems. The advance of network and distributed computing technologies have made it possible for them to collaborate in primitive ways such as and sending messages. There are however great needs to intergrate those isolated information and process islands at a higher level so that they can collaboratively provide business solutions that each individual application is unable to provide.

Workflow has been proposed to address the above problems of early computer applications. The basic idea behind the workflow technology is to separate business process and workflow management component from the existent applications to increase flexibility and maintainability. The major driving forces of workflow are the need for business re-engineering whose purpose is to increase productivity, reduce cost, and respond to the changing environment more quickly, and the advent of technologies such as distributed computing, object technology, databases, etc. that facilitate open and reliable information exchange and collaboration across the organization.

The separation of business policies from applications not only makes (workflow-based) applications easier to maintain and enhance, as changes in procedures can be made using workflow tools without having to rewrite the application, but also provides several other advantages. For example, a workflow system can dramatically increase productivity, as business procedures are automated. A workflow system supports policy-driven allocation of resources and can therefore adapt dynamically to changing workloads. Since workflow processes can be understood by computers, it is also possible to develop workflow tools that track process executions and control process execution in a more flexible way. Another big advantage of workflow systems is that it simplifies application development, not only because application components can be reused, but also because functions common to many applications such as recovery have already been provided (by the underlying workflow management systems).

1.2 Workflow Systems

Workflow management is a diverse and rich technology and is now being applied over an ever increasing number of industries. Workflow is also a generic term which may refer to different things at different levels such as process modeling at the business process level, or process specification and enactment at the system level. In this paper, we discuss issues on process specification and enactment for various kinds of workflow systems (e.g., ad hoc, administrative, production, and collaborative workflow systems). The business perspective, for example, issues such as business re-engineering, process modeling, and BPR tools, will not be covered.

A workflow process, as defined in [WfMC94b], is a coordinated (parallel and/or serial) set of process activities that are connected in order to achieve a common business goal. A process activity is defined as a logical step or description of a piece of work that contributes toward

the accomplishment of a process. A process activity may be a manual process activity and/or an automated process activity. A workflow process is first specified using a process definition language and then executed by a workflow management system (WFMS). A WFMS is a system that completely defines, manages and executes workflow process through the execution of software whose order of execution is driven by a computer representation of the workflow process logic.

It has been over ten years since the first workflow product was introduced. There are now at least several dozens of workflow products available on the market with certain workflow capabilities. Workflow technology has been used in a wide range of application areas such as banking (loan application), finance (stock), insurance (claim process), health care, telecommunication, manufacturing, and document management.

Despite all these efforts, workflow is far from pervasive. There are many reasons for the situation. The following are some of the major ones, from a technical point of view,

Infrastructure. Workflow systems are much more than just workflow engines that execute workflow processes. Successful execution of a workflow process requires proper support from the underlying infrastructure. For example, technologies such as distributed computing, object orientation, and security are necessary for the workflow engine to invoke external applications (especially legacy applications). Unfortunately, distributed computing and object technologies such as CORBA and ActiveX/DCOM have not been mature enough for real applications until recently.

Standards. The lack of standards has been one of the major obstacles to wide application of workflow technology. Unlike relational databases, each workflow vendor has its own workflow model, specification language, and API. Recent efforts by the Workflow Management Coalition (WfMC) have made significant progress, but there is still a long way to go.

Complexity. Workflow application development is a complex task involving more than simply specifying a process definition, which itself is a formidable task. Other (and harder) tasks include wrapping external applications to be invoked by the workflow engine, managing workflow resources, and setting up communication infrastructure. Unfortunately, current workflow systems provide little help for facilitating these tasks. Every major workflow applications require lengthy and intensive collaboration between the workflow vendors and the application developers.

Technology. Despite all the technical progress, workflow technology is still far from mature. It is evident that none of the existing workflow products or research prototypes can provide the same level of support as relational database management systems do for reliable and consistent process execution. It is true that many workflow applications do not need this level of support. But it is also important for the workflow

management system to have the ability so that mission critical applications that are currently implemented using other technologies (e.g., database) can be re-engineered using workflow.

There are other papers discussing the limitations of existing workflow products and outlining important research issues (see, e.g., [AAEM97], [MAG⁺95] and [VLP95]). This paper focuses on technical solutions in order to make workflow more pervasive. We discuss the current capabilities of workflow products and major issues that need to be addressed before workflow can be successful in the market place.

2 Workflow Products

In this section, we first summarize the major features, enabling technologies, and successful applications of the current generation of workflow systems. We then describe a few industry trends that we believe are important to the next generation of workflow systems.

2.1 Current Status

Workflow systems have evolved at least three generations according to [MB91]. The first generation of workflow systems are monolithic applications of a particular application area (e.g., image or document management). The second generation workflow systems factored out the workflow components but they are still tightly coupled with the rest of the products. The third generation workflow systems have generic, open workflow engines which provide an infrastructure for robust production-oriented workflow. The workflow specification is given separately through a graphical user interface and is interpreted by the workflow engines. [VLP95] has also predicted a fourth generation of workflow systems that are part of the middleware and will offer services among other services.

2.1.1 Workflow Product Features

Workflow products of the early age (e.g., WorkFlo by FileNet) are image-based. The purpose of such systems is to automate and manage the flow of images, data, text, and other information throughout an organization. These systems are thus also data- or document-centric. The main function of a data-centric workflow process is to route the data (e.g., a design document) so that people can work on the data.

In recent years, most workflow vendors have either developed or relabeled their products as non image-based. Most of the workflow products are also process-centric (instead of data-centric) in the sense that workflow processes formalize and enforce business policies. On the other hand, there are still needs for data-centric workflow products and there are vendors such as XSoft focusing on the segment.

In the following, we summarize major features of the current generation of workflow products that are non image-based and process-centric. Note that this is not a complete list and the listed features may also not be supported by all workflow products. Nevertheless, we believe that these features characterize the current generation of workflow products and are supported (at least partially) by most workflow products.

Graphical representation. Perhaps the most significant improvement of the current generation of workflow products over earlier generations is the ability to specify and represent workflow processes as graphical maps. In the map, major workflow steps, data and control flows, as well as other components of a workflow process are displayed graphically using icons and lines connecting icons. The idea is to provide an intelligible process view to non-programmers such as business analysts, re-engineering consultants, end-users and supervisors.

The workflow process map is now a standard component of all workflow products. But it differs significantly from vendor to vendor with respect to the information contained in the map and the way it is represented. For example, some products support only a single one-level map while others represent process maps hierarchically, i.e., that there is one main map which contains icons representing submaps. Some products manipulate workflow data explicitly on the map while other do not. The granularity of process maps is also different. For example, some products do not include specifics of the workflow activities. As noted in [Silv95], most workflow products fail to describe the workflow process with sufficient clarity and completeness.

Architecture. Workflow systems, by nature, are distributed systems. Most workflow systems employ client/server architecture and run on multiple platforms. There is a workflow engine which acts as a coordinator. Other components of workflow systems such as process monitor, process starter, and process controller are all clients of the engine. External applications that perform workflow tasks can be both geographically dispersed and on different platforms.

The workflow engines in most existing workflow systems are still centralized in the sense that the entire execution of a process is handled by a single workflow engine (or a cluster of engines that share the same data storage). It is possible in some workflow systems to start a subprocess at a different machine as a step of the containing process execution. But the subprocess execution and the containing process execution can be considered as separate process executions with little interaction except data passing at the beginning and end of the subprocess execution. No workflow systems can support reliable and consistent process execution collectively by more than one independent (share nothing) workflow engine.

Data model. In the WfMC workflow reference model, three kinds of workflow data have been identified: *process control data* that is manipulated by the workflow management system only; *process relevant data* that is used by both the application and workflow

management system; and *application data* that is used by the workflow application only. The idea is to separate business policies (e.g., control flow and data used in flow control) from application details (e.g., data used to perform a task).

All workflow products have their data models but some of them are quite different from the WfMC model. For example, many workflow systems do not distinguish between process relevant data and application data. In these systems, workflow engines have accesses to all workflow data (including application data).

User model. A user model specifies each user's role and the role the user coordinate. The idea is to separate the concept of logical role which is the specification of capabilities needed to perform a task and the concept of physical resources that have certain capabilities to perform the tasks. Process designers specify roles for workflow tasks at design time and specific resources that have the required capabilities will be assigned to the tasks at process execution time. The advantage is the independence of the workflow process from specific resource that may change in the process lifetime.

There are still workflow products that do not distinguish between logical roles and physical resources. However, many workflow products do support this basic user model. Some even support more complicated model that allows specification of a user's organization and manager, the function and processes that the user is authorized to use, etc.

Rule capability. Almost all workflow products allow process executions that are more complicated than simple sequences. Complex flow control requires workflow products to have rule capability. Most workflow products have built-in rule engines. But the rule specification can be quite different. Some products provide script languages for rule specification while others have graphical rule editors that are easier to use.

Tools. One of the advantages of using workflow over monolithic applications is that workflow management systems include tools for process monitoring, tracking, and controlling. Most workflow products provide process development tools and some also provide animation and simulation tools.

2.1.2 Enabling Technologies and Standards

Standards and enabling technologies are important factors in making workflow pervasive. In the past few years, significant progress has been made with respect to workflow related standards such as WfMC, MAPI-WF, and ODBC and enabling technologies such as email, CORBA, and ActiveX/DCOM.

WfMC standards. WfMC was founded in 1993 and is now considered the primary standard body for the workflow market. The standardization work of WfMC is centered around the *workflow reference model* (see Figure 1).

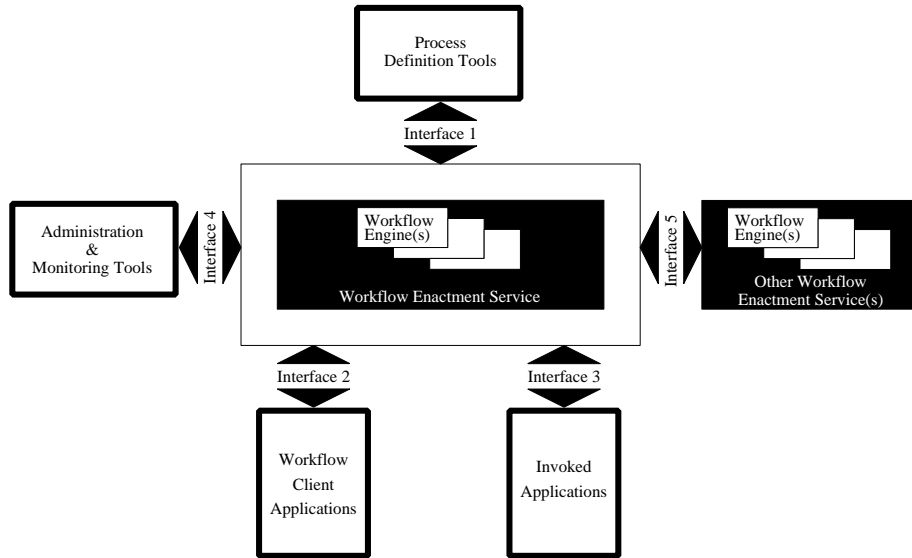


Figure 1: WfMC workflow reference model

The reference model specifies a framework for workflow systems, identifying their characteristics, functions and interfaces. The focus has been on specifying the five APIs that surround the workflow engine. These APIs provide a standard means of communication between workflow engines and clients (including other workflow components such as process definition and monitoring tools). So far, WfMC has draft specifications for all APIs except interface 3. Most workflow vendors plan to support the WfMC APIs and some vendors have already demonstrated the WfMC APIs (e.g., for interface 2) working with their workflow engines.

Workflow interoperability and standards are vital as automation technology becomes more complex and the Coalition's work in this industry is central to keeping up with the rapid progress. On the other hand, workflow standardization is still in its preliminary stage and has a long way to go.

MAPI workflow Framework. MAPI is a message API standard promoted by Microsoft and the MAPI workflow framework (MAPI-WF) is Microsoft's initiative to the WfMC. The idea is to combine the functionalities of workflow systems and the flexibility of messaging systems so that applications that span both messaging users and line-of-business applications can be deployed. It addresses the interoperability issue between messaging systems and workflow systems. In a message environment, a workflow request (e.g., of interface 4) can be packaged within some body part of a message. MAPI-WF provides a standard set of body parts and properties so that workflow packages can be delivered

to and from the workflow engine. Workflow components (e.g., workflow engines, workflow applications, and workflow tools) that conform to MAPI-WF can communicate via messaging systems such as Microsoft Exchange.

Given the popularity of messaging systems and the influence of Microsoft, MAPI-WF will play an important role. Many workflow vendors have already expressed their intentions to support MAPI-WF in their workflow products. To the best of our knowledge, no vendors have actually demonstrated it in their products.

Enabling Technologies. The two most important enabling technologies for workflow systems in recent years are object technology and distributed computing technology. Unlike other software systems such as database management systems, workflow systems are distributed and open by nature. To perform a workflow task, the workflow engine needs to invoke remote workflow applications. Object and distributed computing technologies such as CORBA and ActiveX/DCOM are very useful in wrapping, managing, and invoking heterogeneous applications.

Several workflow products have used CORBA and ActiveX/DCOM as transport services to invoke remote applications. There is also research (e.g., [DKM⁺97]) investigating a CORBA-based workflow enactment system which supports a scalable software architecture, multi-database access, and an error detection and recovery framework.

2.2 Industry Trends

Workflow is a young area and is changing rapidly. Existing workflow products are evolving with new features and new products are being introduced. It is still not very clear what the next generation of workflow products will be. In the following subsections, we list some industry trends that are both important and general enough that they could be adopted by most workflow vendors in the near future. In the next section, we discuss some more advanced issues that are also important but are either not mature enough or not general enough to likely to be widely adopted by workflow vendors in the near future.

2.2.1 Open and Extensible Interfaces

As mentioned, workflow systems are distributed and open by nature. The organizations that will use workflow systems already have computer networks, applications (e.g., spreadsheets), data (stored in files, databases, etc.) and other information. To be useful, a workflow management system must fit into the organization's existing computing environments.

Most existing workflow products include application programming interfaces. External applications can be integrated with the workflow system and external data can be used for workflow process execution via, e.g., application data handlers. To be a really open workflow system, extensible interfaces are needed for incorporating other existing resources and information that are needed in workflow process execution.

For example, events are one of the major means that workflow processes interact with each other and with the external environment. A telecommunication network management process must be able to react to alarms generated by the managed telecommunication network and generate events to effect changes in the network. Although most workflow products still don't support events (especially external events that interact with external environments), there are workflow systems that do. There are also research projects that address the issue.

Another example that requires extensible interfaces is the integration of existing corporate directories. The information about users and the corporate hierarchy is necessary for assigning resources to perform workflow tasks. Requiring users to register themselves to the workflow system as some products do is clearly not the best way. A more flexible way is to provide interfaces to integrate existing corporate directories into the workflow engine. This not only saves workflow application development time, but also makes the later maintenance easier and avoids possible inconsistencies.

There are already workflow products with extensible interfaces for integration of existing resource management systems. More products will provide the interface in near future as this is a feature greatly appreciated by workflow application developers.

2.2.2 Process Development Environments

The development of workflow applications is generally difficult, due to their complexity. The current workflow products address the problem by providing graphical user interfaces for process design and management. GUI tools, however, only address one aspect of the problem which is relatively easier: specifying process templates. A harder problem is to integrate the workflow process with the computing environment. The problem is hard because of the heterogeneity and complexity of the computing environments. To make things worse, most workflow vendors designed and marketed their products as generic tools trying to cover all application areas.

We believe that workflow will not become pervasive until the complexity of developing workflow applications can be significantly reduced. One way of doing that is to provide a good development environment. A good environment must be domain-specific to provide commonly used process templates, commonly used data forms, tools to wrap and manage commonly used applications, basic communication infrastructure, etc. Currently, there are special-purpose workflow products available on the market for specific domains. For example, Araxsys has a workflow product specifically targeting the healthcare market (see [Arax97]).

For the general-purpose workflow products, it is possible to develop special packages based on the general-purpose workflow engine. For example, FileNet has introduced VisualFlo/Payable for account payables. HP has introduced AdminFlow for business administration. More such packages will be expected to cover application domains such as telecommunications, banking, and finance.

As mentioned, an important aspect of a workflow process development environment is to wrap external applications to be used by workflow processes. The wrapped applications can be packaged into a library and then reused by workflow processes. Workflow process development can be further simplified if workflow activities can be reused. Workflow activities include much more than just external applications to be invoked. Other information includes: logical role specification that maps to the specific external application, data needed to perform the task, communication mechanisms, consistency and deadline specification, etc. Unfortunately, most workflow products do not support this level of reuse, as workflow activities contain process-specific information (e.g., position in the process) that cannot be reused. It is thus necessary to separate process-specific and process-independent parts of workflow activities. For example, HP's workflow product distinguishes between the two parts and allows the reuse of the process-independent parts of workflow activities. This allows the special-purpose workflow environment such as AdminFlow to be easily developed based on the general-purpose workflow engine.

2.2.3 Wide Area Workflow

The current generation of workflow products has been criticized for rigid process models, narrow application focus, and platform restrictions (see [Delp97]). These workflow products best service applications where business rules, process flows, and work participants are known in advance and rarely change. The advent of wide area networks and the world wide web has provided new opportunities for workflow. Most workflow vendors have provided web interfaces to their workflow products. There are also research projects trying to develop workflows on the web (see, e.g., [MPSK97]).

As predicted in [Delp97], one possible change for workflow technology is users' environments where workflow tasks are performed. Workflow users will use universal access available via open interfaces such as email, telephone, fax, pager, Web, and intranets/extranets to perform workflow tasks. The key is to separate workflow processes from the user environment so that changes at one side will not affect the other. The major difference between the traditional and wide area workflows is that workflow users have control over what and how they receive information in the latter. The advantage is faster response time and greater productivity by providing multiple access points to the same information and allowing users to use tools of their choice.

3 Workflow Research

Workflow is an active research area with research efforts from both academia and industry (see, e.g., [VLP95] and [MAG⁺95]). On the other hand, workflow research, especially that from academia, has so far made little impact on workflow products. There are two main reasons for the situation. First, early workflow systems, having evolved from different areas such as office automation systems, job control systems, and document management systems,

were struggling to define basic models, architectures, and functionalities, while workflow researchers, most of whom had strong database backgrounds, have been focusing on introducing advanced database techniques to workflow systems. On the other hand, workflow vendors have so far not been very successful in applying workflow technology to applications that require advanced database features such as ACID transactions, and workflow researchers have also failed to develop techniques that are flexible enough for workflow systems.

Nevertheless, we believe that research can address issues that are very important in making workflow pervasive. Our experience with customers shows that there are many workflow applications that require some level of transaction support. The requirements, however, are very different from those in database systems. As a result, not only do existing database techniques need to be adapted to fit into the workflow environment, but also new techniques need to be developed to address issues that are unique in workflow systems.

In the section, we discuss some of the important research issues. We will try to emphasize the differences between the database and workflow environments. Note that this is not a complete list of workflow research issues. The purpose is to inspire research in these and other related areas.

3.1 Transactional Workflow

The concept of transactions was first introduced for database applications in [Gray81]. A transaction is an execution unit with ACID properties: it maps a database from one consistent state to another (*consistency*); either all or none of its effects take place (*atomicity*); and the effects are made permanent once committed (*durability*). Multiple transactions may be executed concurrently, but the overall effect must be equivalent to some sequential execution (*isolation* or *serializability*).

Workflow models that support certain transactional properties have been viewed by many researchers as extensions to the relaxed transaction models (see, e.g., [SR93], [CD96], [GH94], [TV95b], [EL95], [Leym95], [AAE⁺96] and [KR96]). It has been proven both possible and very useful to incorporate transactional semantics such as recovery, atomicity and isolation to ensure correct and reliable workflow executions (see [JLP⁺96]). Database techniques have been adopted to provide transactional properties for workflow processes. For example, failure atomicity is ensured via both forward recovery (see [EL96]) and backward recovery (i.e., compensation, see [Leym95] and [DDSD97]). Execution atomicity can also be ensured by specifying consistency units (or execution atomic units) of workflow processes and coordinating their executions to ensure M-serializability (see, e.g., [BDS⁺93], [RS94] and [TV95b]).

On the other hand, a workflow process is fundamentally different from a database transaction as discussed in [WS97]. First, a workflow environment is more complicated than a database and involves heterogeneous and distributed components. A workflow activity can be very complicated and involve human interactions. Second, a workflow process is structurally

more complex than a database transaction, and the execution of a process may establish quite complex control and data flow dependencies among the activities of the process. A workflow process specification may include conditional branching, concurrent execution of activities, loops, and other complex control structures.

Database recovery techniques such as logging have been successfully adopted in workflow systems. There are workflow products that support reliable workflow process executions. Ensuring atomic and consistent process execution, however, is still missing from workflow products and remains an open research issue. In this subsection, we discuss new issues in workflow compensation and concurrency control as the result of the above differences between the database and workflow environments.

3.1.1 Compensation

Compensation was introduced in [GMS87] to simulate the transactional properties for long-running database applications that would be too expensive to implement as single ACID transactions. The idea was to implement such an application as a *saga* or sequence of ACID transactions so that resources needed only at a particular stage could be released after the corresponding transaction completes. Atomicity was simulated by compensating already completed transactions in reverse order.

In workflow systems, compensation is used to deal with process activity failures. When a process activity instance fails, the workflow management system is responsible for bringing the process execution to a designated *save point*, which is a previous execution step of the process. The save point represents an acceptable intermediate state of process execution and hopefully also a decision point where certain actions can be taken to either fix the problem that caused the failure or choose an alternative execution path to avoid the problem. To roll back workflow process execution, compensation activities will be invoked to undo the effects of the completed activities.

Compensation is more complicated (and thus interesting) in workflow systems than in database systems for two reasons. First, compensation specification (i.e., when, what, and how to compensate) is more difficult, due to the complexity of workflow processes and activities. The compensation activity can be as complicated as (or even more complicated than) the original activity to be compensated. Second, optimization of compensation processes (i.e., what activities don't need compensation) is important. Unlike database transactions which can be compensated and re-executed easily and efficiently, workflow compensation can be very costly. It is therefore important to avoid unnecessary compensation as much as possible.

Existing research on the issue has been focusing on static specification of compensation scopes. For example, [Leym95] discussed an enhancement to IBM FlowMark which allows the process designers to specify (at process design time) spheres of compensation to determine

the scope and extent of compensation in case of activity failures. The failure of an activity may cause the compensation of just the failed activity, the entire containing sphere, or the containing sphere and other dependent spheres. A similar approach has also been proposed in [CD97] for hierarchical workflow processes. The compensation scope is determined in a bottom-up fashion: first to the designated save point in the transaction containing the failed activity, then to the designated save point of the higher level containing the transaction if the current level transaction can not handle the failure.

An interesting issue is how to make use of runtime information to further avoid unnecessary compensation. As we mentioned, the purpose of compensation is to undo the effects that caused the failure so that the execution can resume. Thus, a workflow activity needs compensation if it contributed to the failure, and/or its re-execution is different from the original execution. Compensation scopes specify the activities that *might* affect the failed activities in some execution environments. It is however possible that an activity in a statically specified compensation scope did not contribute to a particular failure. For example, an activity a_1 affects a subsequent activity a_2 if another concurrent activity a_3 happened before a_1 . There is no need to compensate a_1 if a_2 failed before a_3 has even started. This information, however, will only be available at run time. Identifying unnecessary compensation and avoiding it at run time is difficult because other nodes may be affected. But it is worth the effort in many cases where compensation and re-execution of workflow activities are very expensive. [DDSD97] presented some preliminary results along this line of research.

In general, we assume some kind of static relationship between the original execution and its compensation. For example, the same compensation strategy will be used for a workflow activity independent of the cause of failures. A compensation activity is defined for each activity or a group of activities that need compensation. This however may not be true in real life. There can be many different ways to recover a failed execution according to the cause of the failure and the compensation process can be structurally independent of the original execution. For example, the compensation to a person who committed a crime is to go through the justice procedure, which has no structure relationship to the original crime procedure. Little research has been done in the area.

The most fundamental issue of compensation is, perhaps, the correct criteria for workflow process execution and compensation. In database systems, a compensation is correct if everything between the save point and the failure point is compensated and in the exact reverse order of the original execution. In workflow systems, we need a more relaxed criterion for optimization purposes. For example, the order requirement could be relaxed for compensations that are commutable. A good understanding of correct compensation is essential to efficient workflow compensation and can even be application-dependent.

3.1.2 Concurrency Control

Concurrency control is a classical technique in databases, which ensures execution isolation of a transaction from other conflicting transactions. Although concurrency control has been considered either unnecessary or too costly for many workflow applications, it can be very important for some workflow applications where mission-critical operation requires a consistent view of the execution environment (see, e.g., [JLP⁺96]).

The problem of concurrency control in workflow systems is however a little bit different from that in database systems. The purpose of concurrency control in database systems is to ensure execution isolation of database transactions which consist entirely of atomic read/write operations that are visible to the DBMS. In workflow systems, the WFMS ensures the execution isolation of workflow activities which consist of atomic read/write operations as well as external executions that are invisible to the WFMS. The WFMS is responsible for the consistency of the overall execution environment which includes both the internal database which is visible to the WFMS and external systems which are invisible to the WFMS, as well as their cross consistency.

The fundamental issue of concurrency control in workflow systems is correctness criteria. Serializability, as used for database transactions, is too strict for most workflow applications. The main reason for this is that workflow activities are generally long-duration. It is unacceptable in many workflow applications to schedule conflicting activities sequentially as for read/write operations in database transactions. Relaxed correctness criteria (which might be application domain-specific) are essential in specifying and enforcing the correct workflow process executions. [KR96] discussed some of the existing research on the subject.

Some existing research addresses the problem by specifying and enforcing data and execution dependencies among workflow activities (see, e.g., [AARS93], [TV95a] and [GH94]). There is also research that adopts database techniques but allows flexible specification of consistency requirements with respect to scope and granularity. For example, [TV95a] allows grouping a collection of workflow activities of a workflow process into a consistency unit and uses traditional concurrency control to ensure isolation of consistency units (in terms of serializability). Correct execution of activities inside a consistency unit is ensured by enforcing proper data and execution dependencies.

3.2 Distributed Workflow Execution

Workflow systems, by nature, are distributed systems. First, external applications that perform workflow tasks are often geographically dispersed. The workflow management system itself can also be distributed. The most common form of distributed WFMS is function distribution. In such a system, different workflow components that perform various workflow functions such as process definition, process execution, process monitoring, and resource assignment run at different sites. WFMS components interact with each other via messages

or remote procedure calls. Another form of distribution is to perform a workflow function by multiple functionally equivalent WFMS components that share common storage (see [KAGM96]). For example, the execution of a workflow process can be collectively performed by several workflow engines sharing the same data storage for process definitions and execution states. Such a system provides better scalability and is resilient to workflow engine failure, but is still vulnerable to data storage failure.

The most difficult form of distribution is for multiple independent WFMSs (sharing no common data storage) to collectively execute a workflow process. In such a distributed system, each WFMS is itself a complete workflow system with its own engine and data storage. There is no centralized server keeping all information about a process execution. Such a system may be preferred for performance or reliability reasons. The system is more efficient because workflow activities can be executed by the WFMSs that are close to the corresponding external applications (thus reducing communication cost between the WFMSs and applications) and because the WFMSs access process definitions and execution states locally (thus reducing communication cost between the WFMSs and the data storage). It is also more reliable because the failure of one or more WFMSs (including the corresponding data storage) does not stop workflow process executions. The overall system is functional as long as one of the WFMSs is still running.

There are two main issues in implementing such a distributed workflow system: data replication and execution coordination. Data replication is necessary to ensure reliable process execution. For example, the process execution can survive a single WFMS failure if the process definitions and execution states are replicated at more than one independent (e.g., primary and backup) WFMSs. Data replication (especially that of process execution states), however, can be very costly. Data replication can be provided by the WFMSs, or the underlying systems. The advantage of the workflow system level replication is flexibility. For example, workflow processes can be executed at different levels of reliability from no replication (efficient but vulnerable to single WFMS failure) to full replication (expensive but resilient to single WFMS failure).

Execution coordination is needed when more than one WFMS collectively execute a workflow process. For example, execution of a workflow activity by one WFMS may cause the entire process execution to be suspended, affecting all other WFMSs. The key is to transfer process information to a site when it is needed and in the right order. Static information such as process definitions can be replicated at all relevant sites, but runtime information such as process instance states has to be transferred at run time from site to site. This can be done by either circulating all information pertaining to a process and its execution across different sites, or pre-compiling the process definition to determine at which sites the different activities are to be executed (see [MAG⁺95]). The advantage of the former approach is flexibility in the sense that the WFMS can choose to execute a workflow activity at any site according to the runtime execution environment. The disadvantage is high communication cost as the information package can be very large. The latter approach, on the contrary,

can be efficiently implemented, as only relevant information is transferred to the site, but is inflexible. For example, if a workflow activity is pre-assigned to a site which is not accessible at the time, other sites cannot take over the execution as they do not have the information. Another problem with the approach is that most workflow products assign resources to a workflow activity at runtime. The site that is pre-assigned to execute a workflow activity at process specification time can be far away from the resources (e.g., computer applications) to be invoked.

Concurrency control and compensation may also be complicated when the WFMS is distributed. For example, executions of conflicting activities by different sites have to be coordinated to ensure the consistency of the overall execution. Locking is generally not acceptable as workflow activities are often long running. Serializability, on the other hand, may not be needed for the execution. New correctness criteria and coordination algorithms need to be developed to ensure correct and efficient process execution.

3.3 Dynamic Workflow

One of the common assumptions most workflow research makes is the availability of pre-specified workflow definitions. Although there are research or even workflow products allowing modification of process definition at run time (see, e.g., [CCPP96]), it is still considered to be rare and costly.

Dynamic workflow systems are special workflow systems that have no pre-specified process specifications. They start with some initial activities. When an activity has completed, new activities will be selected according to the execution status and results of the current activity. In other words, the workflow is specified and executed at the same time, which is different from dynamic modifications of pre-specified workflow definitions.

Dynamic workflow systems are suitable for workflow applications where process specifications are frequently modified or cannot be pre-specified. For example, most product designs do not follow strict process. They start with initial tasks (e.g., collecting requirements) and follow the general guidelines. Different tasks are performed in different orders according to the status of the design.

Dynamic workflow requires revisiting most research issues discussed before (as well as issues not mentioned in the paper). For example, specification of consistency requirements will be different, due to the lack of the whole picture of the processes. For the same reason, coordination (especially in distributed environments) of activity executions is difficult. There has been very little research regarding dynamic workflow. Recently, there are some research efforts trying to implement dynamic workflow systems using mobile agents. But such efforts are still in their early stages and have not addressed the aforementioned issues.

4 Conclusions

Workflow management systems, as a solution to address many shortcomings of monolithic computer applications, have attracted interests from both industry and academia. On one hand, there is a high demand for workflow systems for all kinds of computer applications, and on the other hand, workflow is far from pervasive. This paper tries to understand the situation from a technical point of view, focusing on the specification and enactment of workflow processes.

Based on our knowledge and experience, we believe that the following factors have all contributed to the current situation: (1) unavailability of proper infrastructure, (2) lack of standards, (3) complexity of workflow process development, and (4) immaturity of workflow technologies. Despite all the problems, great progress has been made in the last few years with respect to infrastructure, standards, and technologies. In the paper, we discuss both state of the products and state of the art of workflow management systems. The purpose of the paper is to inspire further research and development in some workflow areas that are important or essential to the pervasive of workflow systems.

Acknowledgements

We would like to thank Mary Loomis for her help and support in this writing and for her carefully reviewing the paper which greatly improves the quality and presentation. Thanks also to other members of the OpenPM team, Jim Davis, Ying Huang, and Ming-Chien Shan, for many fruitful discussions.

References

- [AAE⁺96] Alonso, G., Agrawal, D., El Abbadi, A., Kamath, M., Gnthr, R., and Mohan, C., “Advanced Transaction Models in Workflow Contexts”, *Data Engineering*, 1996.
- [AAEM97] Alonso, G., Agrawal, D., El Abbadi, A., and Mohan, C., “Functionalities and Limitations of Current Workflow Management Systems”, *IEEE Expert (Special Issue on Cooperative Information Systems)*, 1997.
- [AGK⁺95] Alonso, G., Gnthr, R., Kamath, M., Agrawal, D., El Abbadi, A., and Mohan, C., “Exotica/FMQM: A Persistent Message-Nased Architecture for Distributed Workflow Management”, *IFIP WG8.1 Working Conference on Information Systems for Decentralized Organization*, 1995.
- [AGK⁺96] Alonso, G., Gnthr, R., Kamath, M., Agrawal, D., El Abbadi, A., and Mohan, C., “Exotica/FMDC: A Workflow Management System for Mobile and Disconnected Clients”, *Distributed and Parallel Databases*, 4(3) 1996.

- [AKG⁺96] Alonso, G., Kamath, M., Agrawal, D., El Abbadi, A., Gnthr, R., and Mohan, C., “Failure Handling in Large Scale Workflow Management Systems”, *IBM Research Report*, RJ9913, 1994.
- [ARM97] Alonso, G., Reinwald, B., and Mohan, C., “Distributed Data Management in Workflow Environments”, *RIDE*, 1997.
- [Arax97] Araxsys, “The Araxsys Solution”, <http://www.araxsys.com>, 1997.
- [AARS93] Attie, P., Aingh, M., Rusinkiewicz, M., and Sheth. A., “Specification and Enforcing Intertask Dependencies”, *VLDB*, 1993.
- [BDS⁺93] Breitbart, Y., Deacon, A., Schek, H., Sheth, A., and Weikum, G., “Merging Application-Centric and Data-Centric Approaches to Support Transaction-Oriented Multi-System Workflows”, *SIGMOD Record*, 22(3), 1993.
- [CCPP97] Casati, F., Ceri, S., Pernici, B., and Pozzi, G., “Semantic Workflow Interoperability”, *EDBT*, 1996
- [CGS97] Ceri, S., Grefen, P., and Sanchez, G., “WIDE: A Distributed Architecture for Workflow Management”, *RIDE*, 1997
- [CD96] Chen, Q. and Dayal, U., “A Transactional Nested Process Management System”, *Data Engineering*, 1996.
- [CD97] Chen, Q. and Dayal, U., “Failure Handling For Transaction Hierarchies”, *Data Engineering*, 1997.
- [DKM⁺97] Das, S., Kochut, K., Miller, J., Sheth, A., and Worah, D., “ORBWork: A Reliable Distributed CORBA-based Workflow Enactment System for METEOR₂”, *The University of Georgia*, UGA-CS-TR-97-001.
- [CR91] Chrysanthis, P. and Ramamritham, K., “A Formalism for Extended Transaction Models”, *VLDB*, 1991.
- [DDS95] Davis, J., Du, W., and Shan, M., “OpenPM: An Enterprise Process Management System”, *IEEE Data Engineering Bulletin*, 1995.
- [DHL91] Dayal, U., Hsu, M., and Ladim, R., “A Transactional Model for Long Running Activities”, *VLDB*, 1991.
- [Delp97] Delphi Consulting Group, “Wide Area Workflow”, *The Delphi Seminars & Institutes*, 1997.
- [DDSD97] Du, W., Davis, J., Shan, M., and Dayal, U., “Flexible Compensation of Workflow Processes”, *Hewlett-Packard Laboratories*, HPL-96-72, 1997.

- [EL95] Eder, J. and Liebhart, W., “The Workflow Activity Model WAMO”, *Coopis*, 1995.
- [EL96] Eder, J. and Liebhart, W., “Workflow Recovery”, *CoopIS*, Brussels, Belgium, 1996.
- [ELLR90] Elmagarmid, A., Leu, Y., Litwin, W., and Rusinkiewicz, M., “A Multidatabase Transactional Model for Interbase”, *VLDB*, 1990.
- [CCPP96] Casati, F., Ceri, S., Pernice, B., and Pozzi, G., “Workflow Evolution”, *ER*, 1996.
- [GMS87] Garcia-Molina, H. and Salem, K., “Sagas”, *SIGMOD*, 1987.
- [GH94] Georgakopoulos, D. and Hornick, M., “A Framework for Enforceable Specification of Extended Transaction Models and Transactional Workflow”, *Journal of Intelligent and Cooperative Information Systems*, 3(3), 1994.
- [Gray81] Gray, J., “The Transaction Concept: Virtues and Limitations”, *VLDB*, 1981.
- [HR87] Haerder, T. and Rothermel, K., “Concepts for Transaction Recovery in Nested Transactions”, *SIGMOD*, 1987.
- [JLP⁺96] Juopperi, J., Lehtola, A., Pihlajamaa, O., Sladek, A., and Veijalainen, J., “Usability of Some Workflow Products in an Inter-organizational Setting”, *IFIP WG8.1 Working Conference on Information Systems for Decentralized Organization*, Norway, 1996.
- [KAGM96] Kamath, M., Alonso, G., Gnthr, R., and Mohan, C., “Providing High Availability in Very Large Workflow Management Systems”, *EDBT*, 1996.
- [KR96] Kamath, M. and Ramamritham, K., “Correctness Issues in Workflow Management”, *Distributed System Engineering*, 3(4), 1996.
- [KYH95] Karlapalem, K., Yeung, H., and Hung, C., “CapBasED-AMS: A Framework for Capability-based and Event-driven Activity Management System”, *CoopIS*, 1995.
- [KTW96] Klingemann, J., Tesch, T., and Walch, J., “Semantics-Based Transaction Management for Cooperative Applications”, *International Workshop on Advanced Transaction Models and Architectures*, India, 1996.
- [KLS94] Korth, H., Levy, E., and Silberschatz, A., “A Formal Approach to Recovery by Compensation Transactions”, *VLDB*, 1990.
- [KS94] Krishnakumar, N. and Sheth, A., “Specification of Workflow with Heterogeneous Tasks in Meteor”, *VLDB*, 1994.

- [Leym95] Leymann, F., “Supporting Business Transactions Via Partial Backward Recovery In Workflow Management Systems”, *BTW*, 1995.
- [LR94] Leymann, F. and Roller, D., “Business Process Management with Flowmark”, *COMPCON*, 1994.
- [MB91] McCarthy, D. and Bluestein, M., “Workflow’s Progress”, *The Computing Strategy Report*, 8(12), Forrester Research, Inc., 1991.
- [MS93] McCarthy, D. and Sarin, S., “Workflow and Transactions in InConcert”, *Data Engineering Bulletin*, 1993.
- [MPSK97] Miller, J., Palaniswami, D., Sheth, A., and Kochut, K., “WebWork: Meteor₂’s Web-based Workflow Management System”, *The University of Georgia*, UGA-CS-97-002, 1997.
- [MAG⁺95] Mohan, C., Alonso, G., Gnthr, R., and Kamath, M., “Exotica: A Research Perspective of Workflow Management Systems”, *Data Engineering Bulletin*, 18(1) 1995.
- [RKT⁺95] Rusinkiewicz, M., Klas, W., Tesch, T., Walch, J., and Muth, P., “Towards a Cooperative Transaction Model - The Cooperative Activity Model”, *VLDB*, 1995.
- [RS94] Rusinkiewicz, M. and Sheth, A., “Specification and Execution of Transactional Workflows”, *Modern Database Systems: The Object Model, Interoperability, and Beyond (W. Kim, Ed.)*, Addison-Wesley, 1994.
- [SDDC96] Shan, M., Davis, J., Du, W., and Chen, Q., “Business Process Flow Management and Its Application in the Telecommunication Management Network”, *Hewlett-Packard Journal*, 47(5) 1996.
- [SR93] Sheth, A. and Rusinkiewicz, M., “On Transactional Workflow”, *IEEE Data Engineering Bulletin*, 1993.
- [Silv95] Silver, M., “BIS Guide to Workflow Software”, *BIS Strategy Decisions*, 1995.
- [TV95a] Tang, J. and Veijalainen, J., “Transaction-oriented Work-flow Concepts in Inter-organization Environments”, *International Conference on Information and Knowledge Management, Baltimore*, 1995.
- [TV95b] Tang, J. and Veijalainen, J., “Enforcing Inter-Task Dependencies in Transactional Workflows”, *CoopIS*, 1995.
- [VLP95] Veijalainen, J., Lehtola, A., and Pihlajamaa, O., “Research Issues in Workflow Systems”, *ERCIM Database Research Group Workshop on Database Issues and Infrastructure in Cooperative Information System*, Norway, 1995.

- [WR92] Wachter, H. and Reuter, A., “The ConTract Model”, *Elmagarmid, A. (ed) Transaction Model for Advanced Database Applications*, Morgan-Kaufmann, 1992.
- [WfMC94a] Workflow Management Coalition, “Workflow Reference Model”, *Workflow Management Coalition Standard*, WfMC-TC-1003, 1994.
- [WfMC94b] Workflow Management Coalition, “Glossary: A Workflow Management Coalition Specification”, *Workflow Management Coalition Standard*, WfMC-TC-1011, 1994.
- [WS97] Worah, D. and Sheth, A., “Transactions in Transactional Workflows”, *Advanced Transaction Models and Architectures (S. Jajodia and L. Kerschberg eds.)*, Kluwer Academic Publisher, 1997.