

Exotica: A Research Perspective on Workflow Management Systems[†]

C. Mohan G. Alonso R. Günthör M. Kamath[‡]

IBM Almaden Research Center
650 Harry Road
San Jose, CA 95120, USA
{mohan, gustavo, rgunther, kamath}@almaden.ibm.com

1 Introduction

Workflow is probably one of the most exciting areas of research that has emerged in the past few years. The concepts and ideas have been around in one form or another for a long time: *computer supported cooperative work, form processing, cooperative systems, office automation*, etc. However, only recently the technology and know-how required to implement commercial systems have been available.

In general, workflow management systems (WFMS) are used to coordinate and streamline *business processes*. These business processes are represented as *workflows*, i.e., computerized models of the business process [11], which specify the individual *activity* steps, the order and the conditions in which the activities must be executed, the flow of data between activities, the users responsible for the execution of the activities, the tools to use with each activity, etc. A WFMS is the set of tools that allow the design and definition of workflows, their instantiation and controlled execution, and the coordination and integration of heterogeneous applications within the same workflow [11]. Users interact with the WFMS by accessing their individual *worklists*, where they can find the activities for which they are responsible without necessarily being aware of the higher level process to which the activities belong. A crucial point to understand workflow management systems is their dependency on a variety of technologies, from databases to distributed processing. When designing a WFMS, the challenge is to build a feasible and common working environment in which all these technologies are integrated in a flexible and easy to use fashion. This integration aspect is precisely what has made workflow management systems so elusive to date and explains the limited success of earlier attempts to provide support for cooperative work. With the first generation of commercial workflow management systems [10], there are still many integration issues that remain to be solved, but there is no doubt that in the future workflow management systems will be pervasive in large corporations.

In this paper we present the Exotica Research Project, currently in progress at the IBM Almaden Research Center. One of the goals of the project is to bring together industrial trends and research issues in the workflow area. It is for this reason that we have focused on a particular commercial product, *FlowMark*, IBM's workflow product. However, our results are easily generalized to general workflow management systems since FlowMark's

[†]This work is partially supported by funds from IBM Hursley (Networking Software Division) and IBM Vienna (Software Solutions Division). Even though we refer to specific IBM products in this paper, no conclusions should be drawn about future IBM product plans based on this paper's contents. The opinions expressed here are our own.

[‡]Currently visiting IBM Almaden from the Computer Science Department at the University of Massachusetts, Amherst, MA 01003, USA.

model is similar to that proposed by the Workflow Management Coalition [11]. In particular, the paper contains a high-level overview of our research in six specific areas that are not product specific. The list of these areas is not, by any means, exhaustive. There are still many issues that remain open. For reasons of space, the paper does not discuss workflow related concepts [11, 26, 12, 24, 27] or the particular details of the architecture and functioning of FlowMark [19, 18, 14, 15]. The interested reader can find more information in the references.

2 Large Scale Workflow Systems

Before discussing workflow management, it is necessary to put in perspective the applications addressed by such systems. Some of the common goals of a WFMS are to achieve better performance of business processes, better quality, enhanced effectiveness, enterprise wide coordination and monitoring, etc. As we believe it is the case with most workflow management systems, FlowMark addresses the coordination of *large* scale business processes. In this there are certain similarities with databases. Small databases for personal computers have an undeniable value, however interesting research and industrial strength products relate to issues such as query optimization, performance, data caching, data mining, schema evolution, or triggers, which only become interesting in large database scenarios. For small applications, it is possible to obtain similar or even more benefits by using a less sophisticated cooperative tool rather than a full-fledged workflow management system. However, when the size of the application grows, and the number of users, sites, and processes increases, only a workflow management system is suitable for the task.

It remains to be determined how “large” is a large application. To give an idea of the magnitude of the problem, consider some of the studies that have been done with FlowMark. In one instance, the number of business processes to be executed concurrently in a period of a month is 300,000. In another case, the problem is to link together more than 4,000 sites geographically distributed over an entire country and coordinate the work taking place over heterogeneous platforms at those sites. A third example involves more than 100,000 users working concurrently on the completion of processes. With these figures, issues that were not previously considered as related to workflow become key aspects to the success of a commercial system. Failure handling, continuous availability, navigational flexibility, and replication, to name a few, are no longer nice features but crucial components of the overall system. It is within this framework that our research is inscribed.

3 Research Issues in Workflow Management

The Exotica project is currently focused on six major research areas. Each of them reflects a distinctive need for workflow management systems but, in most cases, they are related to each other and successful solutions will have to integrate aspects from all of them. The list of these areas is by no means exhaustive, and there are many other issues that remain open.

3.1 Failure Resilience in Distributed WFMSs

Given its goals, the architecture of a workflow management system is necessarily distributed in the sense that its components will reside in different and heterogeneous machines. Furthermore, given its relevance in the control of the business processes of a corporation, the system must be continuously available. Hence, each component must be able to deal with local and communication failures, and the design must be robust enough to avoid stopping the execution of processes even in the event of failures. Such requirements define a wide range of failure handling capabilities that the system must support to be commercially viable. Part of our research has been to analyze the possible failure scenarios and design methods to handle them [1, 2, 17].

In the first place, the execution of an activity within a process involves several components: the database server where the workflow process related data is stored, the workflow server that determines where the activity

is going to be executed, the client interacting with the user, and the client interacting with the actual application performing the activity. All of these components interchange information before, during and after the execution of the activity. This information should not be lost, since this may imply that the results of the execution are not recorded, or that the command to start executing the activity never reaches its destination. Hence, some form of *coordination protocol* must be used amongst all these components to guarantee that all agree on the status of a particular activity. In general, the database where the data is stored is transaction based. In FlowMark, for instance, the data is stored in ObjectStore, an object-oriented DBMS. ObjectStore does not provide either *hot standby* support or a *prepare-to-commit* interface, which makes it complicated to reach a point in the execution where data is guaranteed not to be lost. Moreover, the applications interacting with the workflow management system may not be aware of it, thereby becoming uncooperative and committing work regardless of the state of the workflow management system. We are studying several approaches to these problems: using a persistent message mechanism, providing stable storage in each of the components, and designing a handshake protocol that ensures that information is almost never lost.

A second aspect of failure resilience is the possibility that a component fails while it is executing part of a process. If the failed component is the database, replication can be used to minimize the impact of the failure, as discussed below. If the failed component is part of the workflow management system, means must be devised to allow other components to reroute their operations to components that are still available. This implies using multiple connections among components. We have designed a new architecture for workflow management systems in which the notion of clusters of servers is used to provide enhanced availability of the system [2]. By using several clusters, each of them attached to a separate database, the impact of failures on the overall performance can be reduced. Since each cluster's database contains the same schema information such as process templates, role and staff assignments, and organization definition, process instances can be run at any cluster. Thus, failures will affect only a subset of the process instances and will not prevent the users from starting new instances. The problem that still remains is how to handle the instances that were running in a failed cluster. This particular issue is addressed below when discussing replication in workflow management systems.

3.2 Compensation and Navigation in Workflow Networks

A workflow network, or a process, is the representation of business processes. Its instantiation for execution is a process instance. Any process instance usually encompasses a wide range of activities and involve several participants in the form of computer programs and individuals. At any given moment in time, there may be several instances of the same process being executed simultaneously, with possibly each one being at a different stage of progress. Since processes are defined in advance, it is not possible to evaluate all possible exceptions and error conditions beforehand. A workflow management system must provide means to cope with such situations. Similar ideas are found in advanced transaction models [6] such as Sagas [8] for backward recovery, i.e., how to undo the effects of certain operations, or in Flex transactions [5, 21, 28] for forward recovery, i.e., how to select different paths of execution. However, in workflow systems there are certain errors that may force the user to abort the execution of a process instance, and therefore to lose work that has to be performed again when the process is restarted.

Forward recovery is guaranteed in FlowMark, in the sense that if there are failures, the process is guaranteed to make progress. However, semantic failures are more difficult to address, unless the designer of the process was able to foresee them and introduce the appropriate checks in the flow of control. For backward recovery, FlowMark is being enhanced to provide a form of compensation that allows the user to specify *spheres of compensation* [20] to determine the scope and extent of compensation in case of failures.

Our current research is focused on how to provide a flexible mechanism to navigate through the flow of control of a process in either direction, i.e., forward progress or compensation, and allowing to switch directions several times as necessary during the execution of a process.

3.3 High Availability through Replication

High availability is a key requirement of a workflow management system. Failures should be transparent to the users and should have minimal impact on the normal functioning of the organization. This can only be achieved through replication of the process instance information. As in databases, we use a primary/backup architecture in which process instances run on a primary server, while all actions are also recorded at the backup so that if the primary were to fail, then the backup server can take over the execution of those same instances. As in databases, replication has a high cost in terms of synchronization and lower throughput. Given the large number of processes involved in the application, replication of every single process instance may not be cost-effective. Hence, we define three priority levels with respect to replication. A *hot stand-by* process is fully replicated, i.e., all changes to its data are performed both in the primary and in the backup databases. A *cold stand-by* process is also replicated, but the backup contains only the messages with the information regarding the different steps taken in the execution of the process. When the backup needs to take over the execution of the process, it must first replay those messages to bring the process state up-to-date. This will delay resuming the execution but it reduces the overhead of maintaining the replicated copies. Finally, a *normal* process is not replicated at all. The only guarantee FlowMark provides for these processes is forward recovery, i.e., when the server comes up again, execution will be resumed from where it was left off.

As in DBMSs, besides the replication scheme, it is also necessary to deal with additional problems such as dynamic configuration of the system, incorporation and exclusion of new servers without interrupting the system, message duplication, and so forth. We have addressed all these issues and proposed a replication mechanism for workflow management systems that greatly improves the availability of the system avoiding excessive overhead [17].

3.4 Mobile Computing

Workflow management systems must coordinate users distributed over a wide geographic area. The basic idea being that the processes to be executed are defined and controlled in a centralized server, while the users can perform parts of these processes at remote clients. Each step of a process can be executed anywhere in the system, but after its completion, the results are communicated to the server which in turn prepares the next steps for execution. This is the most common approach in existing systems, and simplifies many problems such as synchronization, concurrency and monitoring, but forces the users to remain *connected to the server* while performing multiple tasks. An increasingly large set of users may not be satisfied with this mode of operation, given the widespread use of portable computers and home computers for office work.

For these users, the common way of operation is to load data in their laptops or desktops by briefly connecting with a server in the office. Then they disconnect from the server, work on that data and, after a few hours or few days, reconnect to the server to transfer the results of their work. This has been identified as one of the most common modes of computing in the future [16]. Note that this includes not only mobile terminals but also desktops or any other computer type connected to the server only occasionally via a modem. Mobile or nomadic computing greatly expands the scope of an organization's distributed computing infrastructure. However, from the workflow management point of view, it becomes more difficult to coordinate the work of many users. Also, note that while workflow systems are tools for cooperation, portable computers are generally viewed as tools for individual work.

To address this problem, we have proposed a design for disconnected client operation based on FlowMark [3]. Among other constraints, the design had to provide support for disconnected clients with minimal changes to the semantics of a business process and its implementation, effectively allowing users distributed over a wide geographic area and working with heterogeneous resources to cooperate, while preserving their mobility and independence. In particular, we define the semantics of loading work to a mobile, disconnectable computer by introducing the notion of *locked activities* and the *user's commitment* to eventually execute them. Locked activities

stay within the user's machine unless they are unlocked. The feasibility of our approach is proven by outlining the implementation issues for FlowMark. We are currently working on a prototype based on these ideas.

3.5 Distributed Coordination

To further enhance the availability and failure resiliency of a workflow management system, it is possible to design it entirely as a distributed system. This implies that processes are transferred from site to site as their execution progresses, without a centralized server keeping all the relevant information. There are two possible approaches. One is to make a large package that contains all the information pertaining to the process and its execution and circulate this package across the different sites. This approach has several problems, mainly the size of the message and the need to maintain multiple copies of it in the system to allow concurrent activities to take place. The other solution is to precompile the process definition to determine at which sites the different activities are to be executed. Once this has been done, the only information that gets transferred from site to site are the results of previous computations and the process state. Using this approach, failures can be made transparent by rerouting a process to different nodes. It also eliminates the potential bottleneck created by a centralized database and server.

Exotica/FMQM [1] proposes a similar architecture in which persistent storage is replaced by persistent messaging. For persistent messaging, IBM has defined an application programming interface (API) standard called Message Queue Interface (MQI) [13], and a family of products called MQSeries that supports MQI [22]. MQSeries products operate on IBM and non-IBM platforms and they support the architected MQI. Communication takes place through named queues that do not require all participating programs to be available, i.e., up and running, simultaneously. Moreover, MQI is not sensitive to network transport protocol differences. The system we propose, *Exotica/FMQM*, FlowMark on Message Queue Manager, is a distributed workflow system in which a set of autonomous nodes cooperate to complete the execution of a process. Each site functions independently of the rest of the system, the only interactions between nodes being conducted through persistent messages which inform about activity completions. This approach, while enhancing the availability and resilience of the system, has the problem of needing more complex mechanisms to monitor and audit the overall execution since there is no centralized server where this information resides.

3.6 Advanced Transaction Models

The goal of most advanced transaction models is to eliminate the constraints imposed by conventional DBMSs on new applications. The requirements of such new applications are completely different from the traditional data processing applications targeted by standard DBMSs. Hence, conventional DBMSs are unsuitable in these new environments. Several advanced transaction models have been proposed [6] but, to this date, most remain as theoretical constructs which have not been implemented. A reason for this state of affairs could be that advanced transaction models are ahead of the available technology and their time is yet to come. We believe, however, that the reason why these models are not being implemented has to do more with their inadequacy to operate in real working environments than with the available technology or applications demands. Advanced transaction models are too database-centric, which provides a nice theoretical framework to work with but limits the possibilities and flexibility of the models. Since they tend to remain theoretical models with no implementation, there are a number of important design issues that are generally not discussed in the literature [23].

Paradoxically, workflow systems are tools to support distributed, heterogeneous environments very similar to those targeted by transaction models. However, the models being used are *workflow* models instead of transaction models. Compare the more than 70 vendors who claim to have workflow systems [10] with the almost absolute lack of commercial products supporting advanced transaction models. Workflow systems bear a strong resemblance to advanced transaction models both in their goals and modeling approach, yet they are quite different in that they address a much richer set of requirements than existing transaction models. As part of our research, we

have analyzed the characteristics of workflow models and the notion of business processes by comparing them with existing transaction models [2]. It is possible to show that the semantics of workflow models are, in general, richer than those of advanced transaction models and more apt to be used in commercial products. Hence, workflow models can be used to implement advanced transaction models. For instance, we have used FlowMark to implement Sagas [7] and Flexible Transactions [28]. Our basic goal was to provide synergy between advanced transaction models and workflow models, an interaction from which both sides may benefit. In doing so, we have developed a better understanding of the inherent limitations of the so-called “advanced models” and identified many points for improvement of workflow systems.

Our approach is not another attempt to merge different transaction models into one or to provide a general framework to program advanced transactions [4]. Transactions are used as the accepted currency in the database community, but our goal is to show the expressibility and power of workflow models compared with what is currently in the research literature. From a purely advanced transactions point of view, our proposal differs from previous ones in that we use an existing commercial product that runs across different platforms to implement different transaction models. This allows us to draw conclusions from direct experience implementing such models in real, working environments. We see workflow models as the the next step after advanced transaction models. They complement many aspects of the latter and address an entirely new range of issues (role and staff assignment and resolution, worklist management, interaction with manual activities, etc.) that make them more suitable for building applications. As has been widely recognized, there is a lack of tools to scale from individual transactions to complex applications [25]. One reason being that advanced transaction models will never reach their technological maturity until they are interpreted in a much broader context. As we show in our research, business processes constitute such a context and workflow models are the tools required to support complex applications.

4 Conclusions

The importance of workflow management systems in large corporations cannot be stressed sufficiently. There is a high demand for such systems, but successful solutions will have to integrate many technologies and approaches. The challenge lies in providing systems capable of dealing with very large, heterogeneous, distributed and legacy applications, while providing an acceptable degree of reliability and availability. The research areas briefly discussed in this paper point towards some of the issues that need special attention in the design of workflow system. These issues apply to the vast majority of existing systems, since they are not constraints imposed by a particular architecture but generic demands of the working environment targeted by these systems. In our case, we have focused on FlowMark since this is an IBM product and we have access to its code and developers. However, the solutions we proposed can be easily extended to other systems given that FlowMark’s model and architecture closely resembles the generic model proposed as part of the Workflow Management Coalition standardization effort [11].

Acknowledgements

Part of this work has been done in collaboration with Amr El Abbadi and Divyakant Agrawal, of the University of California at Santa Barbara.

References

- [1] G. Alonso, D. Agrawal, A. El Abbadi, C. Mohan, R. Günthör, M. Kamath, *Exotica/FMQM: A Persistent Message-Based Architecture for Distributed Workflow Management*, **Research Report, RJ 9912**, IBM Almaden Research Center, November 1994.
- [2] G. Alonso, M. Kamath, D. Agrawal, A. El Abbadi, R. Günthör, C. Mohan, *Failure Handling in Large Scale Workflow Management Systems*, **Research Report, RJ 9913**, IBM Almaden Research Center, November 1994.

- [3] G. Alonso, R. Günthör, M. Kamath, D. Agrawal, A. El Abbadi, C. Mohan, *Exotica/FMDC: Handling Disconnected Clients in a Workflow Management System*, **Research Report**, IBM Almaden Research Center, *in preparation*, January 1995.
- [4] A. Biliris, S. Dar, N. Gehani, H.V. Jagadish, K. Ramamritham, *ASSET: A System for Supporting Extended Transactions*, **Proc. 1994 SIGMOD International Conference on Management of Data**, May 1994, pp. 44-54.
- [5] A.K. Elmagarmid, Y. Leu, W. Litwin, M.E. Rusinkiewicz, *A Multidatabase Transaction Model for Interbase*, **Proc. of the 16th VLDB Conference**, August, 1990.
- [6] A.K. Elmagarmid (editor), *Transaction Models for Advanced Database Applications*, Morgan-Kaufmann, 1992.
- [7] H. García-Molina, K. Salem, *Sagas*, **Proc. 1987 SIGMOD International Conference on Management of Data**, May 1987, pp. 249-259.
- [8] H. García-Molina, D. Gawlick, J. Klein, K. Kleissner, K. Salem, *Coordinating Multi-transaction Activities*, **Proceedings IEEE Spring Compton**, 1991.
- [9] H. García-Molina, D. Gawlick, J. Klein, K. Kleissner, K. Salem, *Modeling Long-Running Activities as Nested Sagas*, **Bulletin of the Technical Committee on Data Engineering, IEEE**, vol. 14, no. 1, pp. 18-22, March 1991.
- [10] C. Frye, *Move to Workflow Provokes Business Process Scrutiny*, **Software Magazine**, April 1994, pp. 77-89.
- [11] D. Hollingsworth, *The Workflow Reference Model*, Workflow Management Coalition, TC00-1003, December 1994.
- [12] M. Hsu, *Special Issue on Workflow and Extended Transaction Systems*, **Bulletin of the IEEE Technical Committee on Data Engineering**, vol. 16, no. 2, June 1993.
- [13] IBM, *Message Queue Interface: Technical Reference*, Document No. SC33-0850-01, April 1993.
- [14] IBM, *FlowMark: Managing Your Workflow*, Document No. SH19-8176-01, September 1994.
- [15] IBM, *FlowMark: Programming Guide*, Document No. SH19-8177-01, September 1994.
- [16] T. Imielinski, B.R. Badrinath, *Mobile Wireless Computing: Solutions and Challenges in Data Management*, **Communications of the ACM**, vol 37, no. 10, October 1994.
- [17] M. Kamath, G. Alonso, R. Günthör, C. Mohan, *Providing High Availability in Very Large Workflow Management Systems*, **Research Report**, IBM Almaden Research Center, *in preparation*, January 1995.
- [18] F. Leymann, W. Altenhuber, *Managing Business Processes as an Information Resource*, **IBM Systems Journal**, vol. 33, no. 2, pp. 326-348, 1994.
- [19] F. Leymann, D. Roller, *Business Processes Management with FlowMark*, **Proc. 39th IEEE Computer Society International Conference (CompCon)**, February 1994, pp. 230-233.
- [20] F. Leymann, *Supporting Business Transactions Via Partial Backward Recovery in Workflow Management Systems*, **Proceedings BTW95**, Dresden, Germany, March 1995.
- [21] S. Mehrotra, R. Rastogi, A. Silberschatz, H.F. Korth, *A Transaction Model for Multidatabase Systems*, **Proceedings of the International Conference on Distributed Computing Systems**, June 1992, pp. 56-63.
- [22] C. Mohan, R. Dievendoff, *Recent Work on Distributed Commit Protocols, and Recoverable Messaging and Queuing*, **Bulletin of the IEEE Technical Committee on Data Engineering**, vol. 17, no. 1, March 1994, pp. 22-28.
- [23] C. Mohan, *Advanced Transaction Models - Survey and Critique*, Tutorial presented at the **ACM SIGMOD International Conference on Management of Data**, 1994.
- [24] B. Reinwald, *Workflow Management in verteilten Systemen*, *Teubner-Texte zur Informatik*, **B.G. Teubner Verlagsgesellschaft**, Stuttgart, Leipzig, 1993.
- [25] B. Salzberg, D. Tombroff, *A Programming Tool to Support Long-Running Activities*, **Technical Report NU-CCS-94-10**, College of Computer Science, Northeastern University, Boston, 1994.
- [26] A.P. Sheth, *On Multi-system Applications and Transactional Workflows*, collection of papers from Bellcore, 1994.
- [27] K. Walk, *Workflow Management Concepts, a White Paper*, IBM Vienna Software Development Laboratory, Austria, November 1993.
- [28] A. Zhang, M. Nodine, B. Bhargava, O. Bukhres, *Ensuring Relaxed Atomicity for Flexible Transactions in Multidatabase Systems*, **Proc. 1994 SIGMOD International Conference on Management of Data**, 1994, pp. 67-78.