

Healthcare Enterprise Process Development and Integration

Kemafor Anyanwu, Amit Sheth, Jorge Cardoso, John Miller, Kryz Kochut

LSDIS Lab, Department of Computer Science

University of Georgia

Athens GA 30602-7404, USA

{anyanwu, amit, jam, kochut}@cs.uga.edu, jcardoso@dme.uma.pt

Tel: 706-542-2310/2911, Fax: 706-542-2966

Healthcare enterprises involve complex processes that span diverse groups and organisations. These processes involve clinical and administrative tasks, large volumes of data, and large numbers of patients and personnel. The tasks can be performed either by humans or by automated systems. In the latter case, the tasks are supported by a variety of software applications and information systems which are very often heterogeneous, autonomous, and distributed. The development of systems to manage and automate these processes has increasingly played an important role in improving the efficiency of healthcare enterprises.

In this paper we look at four healthcare and medical applications that involve investigative, clinical, and administrative functions. Based on these applications, we derive the requirements for developing enterprise applications that involve the coordination of a variety of tasks performed by humans, information systems, and legacy applications.

Categories and Subject Descriptors: H.4.1. [Information Systems]: Information systems applications, office automation, workflow management.

Keywords: workflow management systems, healthcare processes, business processes.

1. INTRODUCTION

The recent push for healthcare reform has caused healthcare organisations to focus on ways to streamline their processes in order to deliver high quality care while at the same time reducing costs. This has precipitated a review and upgrade of clinical and administrative protocols and the increased use of information systems to improve the efficiency of certain processes. Since processes are fundamental building blocks of an organisation's success, information technologies that focus on process management and improvement are good candidates for helping healthcare organisations fulfill their corporate vision.

In the past two decades, a special interest has been taken in Workflow Management Systems (WfMSS) as a tool to streamline, automate, and re-engineer business processes. There are many workflow products which adequately support relatively simple processes, such as document management, form processing, and imaging. However, they fall short in meeting the challenges of

Copyright© 2003, Australian Computer Society Inc. General permission to republish, but not for profit, all or part of this material is granted, provided that the JRPIT copyright notice is given and that reference is made to the publication, to its date of issue, and to the fact that reprinting privileges were granted by permission of the Australian Computer Society Inc.

Manuscript received: 6 September 2002

Communicating Editor: Associate Professor Jim Warren

mission-critical processes, which are often complex, dynamic, large-scale, and QoS-based (Sheth *et al.*, 1996; Cardoso *et al.*, 2002). These qualities are typical in healthcare processes.

Healthcare processes are very complex, involving both clinical and administrative tasks, large volumes of data, and a large number of patients and personnel. For example, an out-patient clinic visit involves administrative tasks performed by an assistant and clinical tasks performed by a doctor or by a nurse. For an in-patient hospital visit, this scenario involves more activities, and the process entails a duration that lasts at least as long as the duration of patient hospitalisation.

Healthcare processes are also very dynamic. As processes are instantiated, changes in healthcare treatments, drugs, and protocols may invalidate running instances, requiring reparative actions (Berry and Myers, 1998; Shrivastava and Wheater, 1998). For example, a care pathway for a patient with disease condition 'A' may need to be changed as new drugs are discovered.

Large-scale processes often span multiple healthcare organisations and run over long periods of time (Dayal *et al.*, 1991). This type of process requires highly scalable workflow systems to support large instances (Bonner *et al.*, 1996). Furthermore, these large-scale processes often need to be integrated with legacy information systems and with distributed, autonomous, and heterogeneous computing environments (Georgakopoulos *et al.*, 1995); thus, they require support for transactional features and error handling (Worah *et al.*, 1997).

Another important requirement is the management of Quality of Service (Cardoso 2002; Cardoso *et al.*, 2002). Healthcare organisations operating in modern markets require Quality of Service (QoS) management. Services with well-defined specifications must be available to patients. An appropriate control of quality leads to the creation of quality care services; these, in turn, fulfill patient satisfaction.

This paper discusses the use of the METEOR workflow system for managing mission-critical healthcare processes. The workflow management and enterprise application integration techniques developed in the METEOR system are intended to reliably support complex, dynamic, large-scale, and QoS-based workflow applications in real-world, multi-enterprise, and heterogeneous computing environments.

An important aspect of the METEOR project is that technology and system development efforts occurred in close collaboration with industry partners. Key healthcare partners have included the Connecticut Healthcare Research and Education Foundation (CHREF), the Medical College of Georgia (MCG), and the Advanced Technology Institute. These collaborations have generated a detailed study of healthcare workflow application requirements, the prototyping of significant healthcare workflow applications with a follow-on trial, and the evaluation of METEOR's technology.

This paper is structured as follows. In Section 2, we discuss the current generation of information systems to support healthcare processes, and we highlight some shortcomings of these systems. Section 3 describes the METEOR system, and Section 4 discusses four healthcare workflow applications that use the METEOR system to meet requirements. Section 5 summarises the benefits of the METEOR approach. Finally, Section 6 presents our conclusions.

2. SUPPORTING HEALTHCARE PROCESSES WITH THE CURRENT GENERATION OF WORKFLOW SYSTEMS

Traditionally, healthcare processes have been managed using limited forms of workflow. Some examples of these are clinical and administrative protocols. However, these "protocols have remained limited in their usefulness in part because developers have rarely incorporated both clinical and administrative activities into one comprehensive care protocol. This lack of integration hinders the delivery of care, as the effectiveness of protocols is often dependent on many administrative tasks being properly executed at the correct time" (Chaiken, 1997). Consequently,

many healthcare organisations are now turning to workflow management techniques to help improve the efficiency of their work processes.

The trend toward computerising business processes has led to a large number of commercially available information systems, some of which specifically target the healthcare sector. These systems offer various levels of process support, functionality, and robustness.

At one end of the spectrum, we have customised workflow application systems that support human-oriented and vertical group processes. These processes typically involve a relatively few number of tasks which are executed in a predefined sequence and which require few roles in a single group of an organisation. In these types of applications, the process model is embedded in the application, and customers need to configure the application in order to tailor it to their specific process.

Some examples include VMI Medical (2002), which offers a pediatric cardiology workflow system; TeleTracking (2002), which enables hospital administrators and staff to effectively manage, coordinate, and deliver quality care to patients; and the Soarian (2002) system, which synchronizes workflows across the entire enterprise and orchestrates patient care by bringing together clinical, financial, therapeutic, and diagnostic information.

Another class of applications at this end of the spectrum focuses on supporting information and document management functions. These applications are usually built on the top of data management systems which are designed to capture, store, retrieve, and manage unstructured information objects such as text, spreadsheets, audio clips, images, video, files, and multimedia.

Some examples include CareFlowNet (2002), which provides for the creation, management, and delivery of medical documentation, and SoftMed (2002), which provides a suite of applications for clinical data management, patient information management, and document acquisition and storage.

At the other end of the process support spectrum, we have workflow management systems which are more general purpose systems. These systems provide tools for process definition, workflow enactment, administration, and for the monitoring of workflow processes.

Research prototypes include METEOR (Kochut *et al*, 1999), MOBILE (Jablonski, 1994), ADEPT (Reichert and Dadam, 1998), EXOTICA (Mohan *et al*, 1995), and MENTOR (Wodtke *et al*, 1996). Commercial products include MQSeries Workflow (IBM, 2002), Staffware (Staffware, 2002), TIBCO InConcert (TIBCO, 2002), and COSA Workflow (COSA, 2002). General information on workflow systems can be found at the Workflow and Reengineering International Association (WARIA, 2002) and the Workflow Management Coalition (WfMC, 2002) Web sites.

The current generation of workflow systems adequately supports administrative and production workflows (McCready, 1992), but they are less adequate for some of the more horizontal healthcare processes which have more complex requirements. These types of processes are dynamic and involve different types of tasks; these can be human-oriented, associated with legacy applications, or associated with database transactions. The processes are large-scale, cross-functional, and cross-organisational, where the different participating groups have distributed and heterogeneous computing environments. Workflow infrastructures to support such processes are limited. This is mainly because many systems have a centralised client/server architecture and support only static processes. They also lack support for features such as exception modelling and handling and QoS management.

Another very important requirement that current workflow systems seldom provide is an integration environment. It is clear that the different functional groups of a healthcare organisation may require different types of applications to support their processes. For example, integrating Picture Archiving and Communication Systems (PACS) with hospital or radiology information systems will allow radiologists to be presented with collateral patient information. This allows for

patient history, clinical information, symptoms, and the previous examination history to be presented to the physician along with images retrieved from the PACS, greatly aiding in the interpretation of images (DeJesus, 1998).

Healthcare organisations typically have various information systems, including legacy applications, that applications that are used routinely and need to be integrated. Unfortunately, many workflow systems of the current generation are based on closed, proprietary architectures. This makes supporting interoperability and integration a complicated, if not impossible, task.

The METEOR system was specifically developed to provide a solution to the problems developed in context of the requirements outlined previously. It supplies an infrastructure that supports mission-critical enterprise-wide processes and that integrates heterogeneous, autonomous, and distributed information systems. A general description of the system is given in the next section. For a comprehensive and detailed description, the reader is referred to Miller *et al* (1998) and Kochut *et al* (1999). Integral to our research was extensive collaboration with our healthcare industry partners. These collaborations resulted in prototyping and deploying several healthcare applications, including a clinical study that is being reported in the Journal of Clinical Pediatrics (Boyd *et al*, 2003).

3. THE METEOR SYSTEM

The METEOR (Managing End to End Operations) system leverages Java, CORBA, and Web technologies to provide support for the development of enterprise applications that require workflow management and application integration. It enables the development of complex workflow applications involving the integration of legacy information systems from multiple organisations with geographically distributed and heterogeneous hardware and software environments. It also provides support for dynamic workflow processes, error and exception handling, recovery, and QoS management. The METEOR system has been successfully used to prototype and deploy several healthcare applications. Our success is due in part to extensive collaboration with our healthcare industry partners.

The METEOR system includes all of the components needed to design, build, deploy, run, and monitor workflow applications. It provides the four main services shown in Figure 1: the Builder, the Enactment, the Repository, and the Manager services.




Figure 1: METEOR system architecture

3.1 The Builder Service

The builder service supports the graphical design of workflows (Lin, 1997; Zheng, 1997). It includes three main components. The task design component provides interfaces with external task development tools, such as Microsoft's FrontPage to design the interface of a user task. The network design component is used to set dependencies, data objects, and transition functions among tasks; it is also used to define security domains and roles. The data design component allows the user to specify data objects that are employed in the workflow.

The service outputs an XML-based representation of process definitions which may be formatted to be compliant with the Workflow Process Definition Language (WPDL) of the Workflow Management Coalition (WfMC, 2002).

3.2 The Enactment Service

There are two enactment services provided by METEOR – ORBWork (Kochut *et al*, 1999) and WebWork (Miller *et al*, 1998). Both services use a fully-distributed open architecture. WebWork is a comparatively light-weight implementation that is well-suited for traditional workflows' help-desk and data exchange applications. ORBWork is better suited for more demanding, mission-critical enterprise applications which require high scalability, robustness, exception-handling support, QoS management, and dynamic modifications.

3.3 The Repository Service

The repository service maintains information about workflow definitions and associated workflow applications. The builder service tools communicate with the repository service to retrieve, update, and store workflow definitions, thereby providing support for rapid application development in the builder service. The builder service tools are capable of browsing the contents of the repository and incorporating fragments (either sub-workflows or individual tasks) of existing workflow definitions into the one currently being created.

A detailed description of the first design and implementation of this service is presented in Yong (1998), and a XML based implementation is described in Arpinar *et al* (2001).

3.4 Management Services

The tools provided by these services are used for administering and monitoring workflow instances. The administration service is used by the workflow administrator to perform management functions, such as installing and configuring workflow instances, load-balancing, and modifying workflow processes in execution. The monitor provides a tool for querying and viewing the state of workflow instances.

3.5 METEOR's Advanced Features

Automatic Code Generation

METEOR has a suitable code generator (Miller *et al*, 1998) that is used to build workflow applications from the workflow specifications generated by the builder service or from those stored in the repository. The code automatically generated from the workflow design stage greatly minimises the steps required to implement the workflow. This frees the designer from having to worry about details of communication or about data passing among existing tasks.

Fully Distributed System

The fully distributed architecture of METEOR yields significant benefits in the area of scalability. METEOR's architecture has three major advantages. First, it allows for the support of workflow

processes that are geographically distributed. Secondly, it provides support for load-balancing among all the participating host machines. Finally, it eliminates the existence of a single point of failure within the system.

The Use of Standards

The METEOR system closely follows the specification and interoperability standards set by bodies such as the WfMC (WfMC, 2002) and the Object Management Group (OMG, 1998). METEOR also supports workflow interoperability standards such as JFLOW (JFLOW, 1998) and SWAP (Swenson, 1998), and it utilises CORBA¹ due to its emergence as an infrastructure of choice for developing distributed object-based applications.

Security

METEOR provides various levels of security, from role-based access control and authentication, to multilevel security (MLS). A MLS workflow system enables globally distributed users and applications to cooperate across classification levels in order to achieve mission-critical goals. Users can program multilevel mission logic to securely coordinate distributed tasks and to monitor the progress of workflows across classification levels (Kang *et al*, 1999).

Dynamic Changes

The METEOR system has a layer that permits consistent realisation of the dynamic change of instances (Chen, 2000). The module guarantees that all consistency constraints which have been ensured prior to a dynamic change are also ensured after the workflow instances have been modified (Reichert and Dadam, 1998).

The features designed to handle dynamic changes in workflows are also very useful in supporting scalability, as the load increases. For example, an administrator may decide to move a portion of a running workflow to a new host (or hosts) that have become available for use.

Error and Exception Handling

Error and exception handling, and the recovery framework (Luo, 2000; Worah *et al*, 1997) have been defined in a scalable manner. The most advanced component developed was the exception-handling mechanism, which works in the following way. During a workflow execution, if an exception occurs, it is propagated to the case-based reasoning exception handling module; the CBR system is used to derive an acceptable exception handler (Luo *et al*, 1998). The system has the ability to adapt itself over time, based on knowledge acquired about past experiences which help solve new problems. As the CBR system collects more and more cases, the global WfMS becomes more and more resistant, thus preventing unwanted states.

QoS Management

The METEOR system allows for the specification of quality of service metrics and requirements (Cardoso, 2002; Cardoso *et al*, 2002). The implementation of mechanisms to specify workflow quality of service (QoS) is a major advance for METEOR. The system includes a workflow QoS model, estimation algorithms and methods, and monitoring tools. The model allows suppliers to specify the duration, quality, cost, and fidelity of the services and products to be delivered. The available algorithms estimate the quality of service of a workflow, both before instances are started

¹ A new version of ORBWork that uses RMI instead of CORBA is currently being implemented.

and during instance execution. The estimation of QoS before instantiation allows suppliers to ensure that the workflow processes to be executed will indeed exhibit the quality of service requested by customers. The analysis of workflow QoS during instance execution allows workflow systems to constantly compute QoS metrics and register any deviations from the initial requirements.

4. HEALTHCARE APPLICATIONS PROTOTYPED USING METEOR

The healthcare sector, including both hospital and non-hospital based organisations (e.g. pharmaceutical companies and laboratories), has a number of different types of organisations. All these organisations have different requirements. Table 1 gives a summary of the different types of processes, the applications that support them, and their requirements.

| | Processes | Example Applications | Requirements |
|---------------------------|---|--|--|
| Hospital Based | Clinical | Charting, Scheduling, Discharge Summaries, Reports | Integration with patient data management software; Management of human and automated activities; Exception handling; Ease of use; Support for Dynamic Changes; Security; Role-Based Authorisation; QoS management. |
| | Non-Clinical (Administrative and Financial) | Ordering Systems (radiology, pharmacy) Patient Management (billing, accounts receivable, claims filing) | Data Management and Integration; Application Integration; Support for Heterogeneous and Distributed Environments; Security; Support for standards (eg. EDI and HL7); Exception Handling |
| Non-Hospital Based | Laboratory | Laboratory Information Systems | Scalability; Exception Handling; Management of complex data types; Transactional Workflows; Integration with other systems; Support for HAD environments; QoS management |
| | Pharmaceutical Industry | Clinical Drug Trial Management | Distributed Environment; Scalability; Exception Handling; QoS management |

Table 1: Healthcare Processes and Applications

The rest of this section describes four out of the six healthcare applications that we have prototyped using the METEOR system. These applications support different types of processes,

4.2 GeneFlow

GeneFlow was developed specifically for the needs of the Fungal Genome Initiative. This is a multi-institution consortium of research groups which is mapping and sequencing the genomes of important fungal organisms.



Figure 3: Workflow design for GeneFlow

GeneFlow is a workflow application that handles the needs of data analysis for genome sequencing. Raw “shotgun” DNA sequence data consists of short overlapping DNA sequences. This data comes from automatic sequencing machines. From this raw data, the short overlapping shotgun sequences must be synthesised into larger contiguous sequences of whole chromosomes. These larger sequences are searched for probable genes and other chromosomal features. The results are then electronically published, with the objective of making the annotated genomes available in the public domain.

Genomic projects involve highly specialised personnel and researchers, sophisticated equipment, and specialised computations involving large amounts of data. The characteristics of the human and technological resources involved, often geographically distributed, require a sophisticated coordination infrastructure to manage not only laboratory personnel and equipment, but also the flow of data generated.

The quality of service management is an important factor for this application (Cardoso, 2002). The laboratory wishes to be able to state a detailed list of requirements for the service to be rendered to its customers. As an example, requirements may include the following constraints:

- The final report has to be delivered in 31 weeks or less, as specified by the customer (e.g. NIH).
- The profit margin has to be 10%. For example, if a customer pays \$1,100 for a sequencing, then the execution of the GeneFlow workflow must have a cost for the laboratory that is less than \$1,000.
- The error rate of the task *Prepare Clones and Sequence* has to be at most ϵ , and the data quality of the task *Sequence Processing* has to be at least α .

- In some situations, the client may require an urgent execution of sequencing. Therefore, the workflow has to exhibit high levels of reliability, since workflow failures would delay the sequencing process.

In this application, METEOR tools are used to wrap genome data analysis applications together in a “genome data assembly line.” Three heterogeneous platforms (SGI, Solaris, and Windows) are used with a single database and a single workflow system. The process requires many human and automated tasks, support for legacy applications integration, and Web-based access to support geographically distributed users. The integration of legacy applications on the SGI, Solaris, and Windows platforms was accomplished by writing Java wrappers for the legacy tasks. These wrappers were then easily integrated with the ORBWork enactment service.

The genetic workflow application presented underlines QoS management requirements. It necessary to analyse the QoS of workflows during the design phase and also during the execution of instances. At runtime, the system monitors instances and registers any deviations from the initial requirements. When deviations occur, the dynamic change interface can be used to adapt workflow instances, with the goal of restoring their QoS to acceptable metrics.

4.3 Eligibility Referral

The Eligibility Referral application was developed for the Connecticut Healthcare Research and Education Foundation (CHREF) to support the process of transferring a patient from one hospital to another. It involves three organisations, two hospitals, and an insurance company.

The design depicted in Figure 4 shows a consolidated workflow, including the activities carried out by both the sending and the receiving hospitals.

The workflow starts with the sending hospital trying to determine the right placement for a patient that needs to be sent out. Once this is done, the next tasks involve determining the eligibility information, obtaining the necessary payment information, and also getting the physician’s signature for a specific patient. The final step in the sending hospital’s workflow is to receive an acknowledgment from the receiving hospital indicating that it will accept the patient. Once this is done, the sending hospital can update its database, and the receiving hospital will take over from there. The



Figure 4: Eligibility Referral Workflow

receiving hospital also has its own workflow for processing transferred patients. Workflow instances spans across the two hospitals, interacting with the insurance company through EDI transactions.

The Eligibility Referral application requires an infrastructure that supports distributed and heterogeneous environments. Workflow instances must be managed across multiple workflow system installations. The application accesses multiple databases and web servers. Furthermore, it requires an infrastructure that supports heterogeneous tasks such as human, automated, and transactional tasks with EDI transactions.

In our implementation, we have deployed separate METEOR systems – one for the sending hospital and one for the receiving hospital. A single workflow instance executes tasks across both the hospitals. Each hospital hosts its own web server and a database. The databases are used to find data about patients in order to verify eligibility information.

4.4 State-Wide Immunisation Tracking

According to the Health Plan Employer Data and Information Set (HEDIS), the childhood immunisation rate is one of the most important elements that define the quality of care. Consequently, childbirth reporting and immunisation tracking are two important criteria incorporated in the performance monitoring and reporting frameworks used in healthcare management. The Immunisation Tracking application has the most advanced requirements of all four examples discussed. The workflow application spans several organisations, the central location, the Connecticut Healthcare Research and Education Foundation, Inc. (CHREF), health-care providers (Hospitals, Clinics, and home healthcare providers), and user organisations (State Department of Health (SDOH), schools, and Department of Social Services (DSS)). It involves 13 tasks including tasks for the admit clerk, triage nurse, and for the eligibility check.

The schematic in Figure 5 shows the system in terms two subsystems: the Clinical and Tracking subsystems. The Clinical subsystem provides features for managing clinical processes such as



Figure 5: Schematic view of the Immunization Tracking application

Worklist Management for the different clinical roles, Admit Clerk, Triage Nurse, Nurse Practitioner, and Doctor; Automatic generation of Medical Alerts (e.g. delinquent immunisations) and Insurance Eligibility Verification by the Admit Clerk; and Generation of contraindications for patients visiting a hospital or clinic, to caution medical personnel regarding procedures that may be performed on the patient. The Tracking subsystem involves reminding parents and guardians about shots that are due or overdue and informing field workers about children who have not been receiving their immunisations.

The development of the application took into account some specific user requirements such as:

- support for transparent coordination of tasks across a distributed and heterogeneous client/server-based architecture in a heterogeneous computing environment;
- support for security measures that preserve patient confidentiality;
- support for a variety of tasks: (non)/transactional, human, and application;
- capability to use existing infrastructure such as DBMS and standards (e.g. EDI)
- low cost, ease of use, modification (re-design), scalability, extensibility, and rapid prototyping and deployment.

Figure 6 shows the system test-bed for the Immunisation Tracking application. It shows the heterogeneous (Solaris 2.4, Windows/NT, Windows95) and distributed (locations in Georgia and Connecticut) computing environment infrastructure, with multiple Web servers, CORBA servers, and multiple databases (five databases on two DBMS systems: Illustra and Oracle).

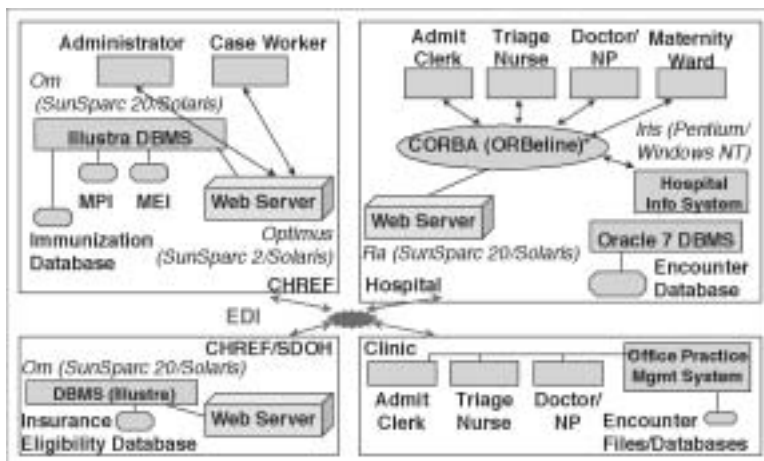


Figure 6: Implementation test-bed for the Immunization Tracking application

5. PROCESS MANAGEMENT CAPABILITIES AND BENEFITS FOR HEALTHCARE APPLICATIONS APPROACH

In this section, we review the promise of recent advances in distributed computing infrastructure, middleware, and Web technologies, exemplified by the METEOR system, to meet the requirements of healthcare applications discussed in Section 4. One important capability is the ability to quickly integrate applications and information systems to support complex and dynamic business process management. Table 2 provides a summary of features, which based on our experience, have been identified as requirements for prototyping healthcare workflow applications.

| Example Application | Capabilities | Benefits |
|---|--|---|
| All | Graphical building of complex applications | Ability to visualise all application components; reduced needs for expert developers; rapid deployment. |
| All except Clinical Pathways | Support for heterogeneous, distributed computing environment; open-systems architecture and use of standards | Seamless deployment over networked heterogeneous (Solaris and NT) server platforms; ease of integration of legacy/existing applications; appeal to customers preferring non-proprietary and multi-vendor solutions. |
| All | Automatic code generation | Significantly reduced coding and corresponding savings in development cost; reduced need for expert developers; rapid deployment. |
| All | Integration of human and automated activities | Natural modeling of complex business activities/processes. |
| All except Clinical Pathways & GeneFlow | Full distributed scheduling | High scalability and performance, minimal single point of failure. |
| GeneFlow | QoS management | Specification, analysis, and monitoring of quality of service metrics. |
| None | Dynamic changes | Rapidly adapt to changes in business processes. |
| All | Traditional security | Support for roles and security on open internetworking. |
| Eligibility Referral & IZT | Database middleware support | Simplified access to heterogeneous variety relational databases on servers and mainframes. |
| Eligibility Referral & IZT | Workflow interoperability standards | Integration with other vendor's products, interoperability in multi-vendor and inter-enterprise applications such as e-commerce. |
| None | Transaction support, exception-handling and automatic recovery, survivability | 7x24 operation and support for mission-critical processes. |
| All | Different levels of security (roles, authorisation, network) | Flexible support for a broad range of security policies. |
| None | Component repository | XML-based reusable application components for rapid development of new applications. |

Table 2: Benefits of the METEOR approach

6. CONCLUSION

Based on the deployment of real-world workflows using the METEOR system, we have drawn a set of requirements for workflow systems supporting healthcare applications. Today's healthcare processes require capabilities for mission-critical workflow support and enterprise integration. Indispensable features include the seamless deployment over networked and heterogeneous server

platforms; rapid deployment of applications; ease of integration of legacy/existing applications; high scalability and performance; specification, analysis, and monitoring of quality of service metrics; and adaptation to changes.

The METEOR system provides a number of features that support the requirements of prototyping and deploying healthcare workflow applications. First, it enables a rapid design-to-development via automatic code generation. Its workflow model and enactment system support a variety of indispensable activities – user and application (automated) tasks – to be used in real-world organisational processes. The workflow engines support heterogeneous and distributed computing environments. This allows workflow process and data to be distributed within and across enterprises. Reliability is an inherent part of the WfMS infrastructure; it includes support for error handling and recovery by exploiting transaction management features. A well-defined hierarchical error-model is used for capturing and defining logical errors, and a recovery framework provides support for the detection and recovery of workflow system components in the event of failure. The system also supports a dynamic change interface, QoS management, and a case-based reasoning subsystem to effectively handle exceptions.

7. ACKNOWLEDGEMENTS

The METEOR team consists of Kemafor Anyanwu, Jorge Cardoso, Prof. Amit Sheth (PI), Prof. Krys Kochut (co-PI), and Prof. John Miller (co-PI). Key past contributors include: Ketan Bhukhanwala, Zhongqiao Li, Zonghwei Luo, Kshitij Shah, Souvik Das, David Lin, Arun Murugan, Devanand Palaniswami, Richard Wang, Devashish Worah, and Ke Zheng.

This research was partially done under a cooperative agreement between the National Institute of Standards and Technology Advanced Technology Program (under the HIIT contract, number 70NANB5H1011) and the Healthcare Open System and Trials, Inc. Consortium. Additional partial support and donations were provided by Iona, Informix, Hewlett-Packard Labs, and Boeing.

REFERENCES

- ARPINAR, I. B., MILLER J. A. and SHETH A. P. (2001): *An efficient data extraction and storage utility for XML documents*. Proceedings of 39th Annual ACM Southeast Conference, Athens, GA. 293–295.
- BERRY, P. M. and MYERS K. L. (1998): *Adaptive process management: An AI perspective*. ACM Conference on Computer Supported Cooperative Work, Seattle, Washington.
- BONNER, A., SHRUF A. and ROZEN S. (1996): *Database requirements for workflow management in a high-throughput genome laboratory*. Proceedings of the NSF Workshop on Workflow and Process Automation in Information Systems: State-of-the-Art and Future Directions, Athens, GA. 119–125.
- BOYD, R., MURDISON, K., BAFFA, J., BRUMUND, M., SHETH, A., KARP, W. and BHATIA, J. (2003): A low-cost web-based tool for pediatric echocardiographic consultations, *Clinical Pediatrics*, 2003 (to appear).
- CARDOSO, J. (2002): Quality of service and semantic composition of workflows. Ph.D. Dissertation. Department of Computer Science, University of Georgia, Athens, GA.
- CARDOSO, J., SHETH A. and MILLER J. (2002): *Workflow quality of service*. International Conference on Enterprise Integration and Modeling Technology and International Enterprise Modeling Conference (ICEIMT/IEMC'02), Valencia, Spain, Kluwer Publishers.
- CAREFLOWNET (2002): CareFlowNet Home Page, <http://www.careflow.com>
- CHAIKEN, B. (1997): Workflow in Healthcare, <http://www.araxsys.com/araxsys/resources/workflow.pdf>
- COSA (2002): COSA Workflow. <http://www.ley.de/cosa/index.htm>
- DAYAL, U., HSU M. and LADIN R. (1991): *A transactional model for long-running activities*. Proceedings of the 17th International Conference on Very Large Databases. 113–122.
- DeJESUS, X. (1998): Integrating PACS Power. *Healthcare Informatics* 15(20): 97.
- GEORGAKOPOULOS, D., HORNICK M. and SHETH A. (1995): An overview of workflow management: From process modeling to infrastructure for automation. *Distributed and Parallel Databases, An International Journal* 3(2): 119–153.
- IBM (2002): MQSeries Workflow. <http://www-3.ibm.com/software/ts/mqseries/workflow/>
- JABLONSKI, S. (1994): *MOBILE: A modular workflow model and architecture*. Proceedings of the 4th International Working Conference on Dynamic Modelling and Information Systems, Noordwijkerhout, Netherlands.

- JFLOW (1998): OMG BODTF RFP #2 Submission, Workflow Management Facility, Revised Submission, <ftp://ftp.omg.org/pub/docs/bom/98-06-07.pdf> 4(7).
- KANG, M. H., FROSCHER, J. N., SHETH, A. P., KOCHUT, K. J. and MILLER J. A. (1999): *A multilevel secure workflow management system*. Proceedings of the 11th Conference on Advanced Information Systems Engineering, Heidelberg, Germany, Springer-Verlag. 271–285.
- KOCHUT, K. J., SHETH, A. P., MILLER, J. A., ARPINAR, I. B. and CARDOSO, J. (2003): IntelliGEN: A distributed workflow system for discovering protein-protein interactions, *Distributed and Parallel Databases*, An International Journal, *Special Issue on Bioinformatics*, 13(1): 43–72
- KOCHUT, K. J., SHETH, A. P. and MILLER, J. A. (1999): ORBWork: A CORBA-based fully distributed, scalable and dynamic workflow enactment service for METEOR, Large Scale Distributed Information Systems Lab, Department of Computer Science, University of Georgia, Athens, GA.
- LIN, C. (1997): A portable graphic workflow designer. M.Sc. Thesis. Department of Computer Science, University of Georgia, Athens, GA.
- MCCREADY, S. (1992): There is more than one kind of workflow software. *Computerworld*. November 2: 86–90.
- MILLER, J. A., PALANISWAMI, D., SHETH, A. P., KOCHUT, K. J. and SINGH H. (1998): WebWork: METEOR2's web-based workflow management system. *Journal of Intelligence Information Management Systems: Integrating Artificial Intelligence and Database Technologies (JIIS)* 10(2): 185–215.
- MOHAN, C., ALONSO, G., GUENTHOER R. and KAMATH M. (1995): Exotica: A research perspective of workflow management systems. *Data Engineering Bulletin* 18(1): 19–26.
- OMG (1998): BODTF RFP #2 Submission, Workflow Management Facility, Revised Submission, <ftp://ftp.omg.org/pub/docs/bom/98-06-07.pdf>
- REICHERT, M. and DADAM P. (1998): ADEPTflex – Supporting dynamic changes of workflows without losing control. *Journal of Intelligent Information Systems – Special Issue on Workflow Management* 10(2): 93–129.
- SHETH, A., GEORGAKOPOULOS, D., JOOSTEN, S., RUSINKIEWICZ, M., SCACCHI, W., WILEDEN J. and WOLF A. (1996): Report from the NSF workshop on workflow and process automation in information systems, Department of Computer Science, University of Georgia, Athens, GA, Technical Report UGA-CS-TR-96-003, October.
- SHRIVASTAVA, S. K. and WHEATER, S. M. (1998): *Architectural support for dynamic reconfiguration of distributed workflow applications*. IEEE Proceedings Software Engineering. 155–162.
- SOARIAN (2002): Soarian Home Page, <http://www.smed.com/solutions/products/soarian/index.php>
- SOFTMED (2002): SoftMed Home Page, <http://www.softmed.com/>
- STAFFWARE (2002): STAFFWARE. <http://www.staffware.com/>
- SWENSON, K. (1998): SWAP – Simple Workflow Access Protocol. Workshop on Internet Scale Event Notification, Irvine, CA.
- TELETRACKING (2002): TeleTracking Home Page, <http://www.teletracking.com/>
- TIBCO (2002): TIBCO InConcert. http://www.tibco.com/products/in_concert/
- VMIMEDICAL (2002): VMImedical Home Page, <http://www.vmimedical.com/>
- WARIA (2002): Workflow and Reengineering International Association, <http://www.waria.com/>
- WfMC (2002): Workflow Management Coalition, <http://www.wfmc.org/>
- WODTKE, D., WEISSENFELS, J., WEIKUM, G. and DITTRICH, A. K. (1996): *The MENTOR project: Steps towards enterprise-wide workflow management*. Proceedings of the International Conference on Data Engineering, New Orleans.
- WORAH, D., SHETH, A., KOCHUT K. and MILLER J. (1997): An error handling framework for the ORBWork workflow enactment service of METEOR, LSDIS Lab. Department of Computer Science, University of Georgia, Athens, GA, Technical Report, June.
- YONG, J. (1998): The repository system of METEOR workflow management system. M.Sc. Thesis. Department of Computer Science, University of Georgia, Athens, GA.
- ZHENG, K. (1997): Designing workflow processes in METEOR2 workflow management system. M.Sc. Thesis. LSDIS Lab, Computer Science Department, University of Georgia, Athens, GA.

BIOGRAPHICAL NOTES

Kemafor Anyanwu received a B.Sc. in Biochemistry (1989) and is currently a Ph.D candidate in the Department of Computer Science, at the University of Georgia. She has been a research assistant member of the Large Scale Distributed Information Systems (LSDIS) Lab at the same institution for several years, where she has been involved in research in workflow management and enterprise integration specifically the issues concerning correctness of workflow specification and dynamic workflows. More recently, her research has focused on issues in developing semantic querying languages and techniques and the Semantic Web.



Kemafor Anyanwu

Amit Sheth (<http://lsdis.cs.uga.edu/~amit>) is a Professor of Computer Science, and the director of the Large Scale Distributed Information Systems (LSDIS) Lab, at the University of Georgia. He is also a CTO/Co-founder of Semagix (<http://www.semagix.com>), which has commercialised a Semantic Web technology based on the research at the LSDIS lab. Until 1994 he served in R&D groups at Bellcore, Unisys, and Honeywell. His research has led to two startups, three significant commercial products, several deployed applications, and two patents. His R&D has been in the areas of area of the Semantic Web and semantic interoperability, Semantic Web Process, digital library, distributed workflow management, federated/multidatabase systems, and parallel and distributed information systems.



Amit Sheth

Jorge Cardoso received a B.A. (1995) in Informatics Engineering and a M.S. (1998) in Information Systems and Technologies from the University of Coimbra (Portugal), and he received a Ph.D. (2002) in Computer Science from the University of Georgia (USA). While at the University of Georgia he was part of the LSDIS Lab., his research concentrated on workflow QoS management and semantic composition of workflows. Jorge Cardoso joined the University of Madeira (Portugal) in March 2003 and currently holds a faculty position at the Mathematics and Engineering department.



Jorge Cardoso

John A. Miller is a Professor of Computer Science at the University of Georgia and is also the Graduate Coordinator for the department. His research interests include database systems, simulation and workflow as well as parallel and distributed systems. Dr. Miller received his the M.S. and Ph.D. in Information and Computer Science from the Georgia Institute of Technology in 1982 and 1986, respectively. He is an Associate Editor for ACM Transactions on Modeling and Computer Simulation and IEEE Transactions on Systems, Man and Cybernetics.



John A. Miller

Krys J. Kochut is a Professor and Head of the Computer Science at the University of Georgia. His research interests include distributed systems (especially workflow systems), database systems, and bioinformatics. He is responsible for design and development ORBWork, the first CORBA and Java based distributed workflow management systems, which has been available as a commercial product. Also, he developed FGDB, a software system for the construction of physical maps of genomes. He has published over fifty technical papers, as well as received several research grants from National Science Foundation, National Institute of Standards and Technology, and Naval Research Laboratory.



Krys Kochut