Wright State University

CORE Scholar

Kno.e.sis Publications

The Ohio Center of Excellence in Knowledge-
Enabled Computing (Kno.e.sis)

1999

Processes Driving the Networked Economy

Amit P. Sheth

Wright State University - Main Campus, amit@sc.edu

Will van der Aalst

I. Budak Arpinar

Follow this and additional works at: <https://corescholar.libraries.wright.edu/knoesis>

Part of the [Bioinformatics Commons](#), [Communication Technology and New Media Commons](#), [Databases and Information Systems Commons](#), [OS and Networks Commons](#), and the [Science and Technology Studies Commons](#)

Repository Citation

Sheth, A. P., van der Aalst, W., & Arpinar, I. B. (1999). Processes Driving the Networked Economy. *IEEE Concurrency*, 7 (3), 18-31.

<https://corescholar.libraries.wright.edu/knoesis/816>

This Article is brought to you for free and open access by the The Ohio Center of Excellence in Knowledge-Enabled Computing (Kno.e.sis) at CORE Scholar. It has been accepted for inclusion in Kno.e.sis Publications by an authorized administrator of CORE Scholar. For more information, please contact library-corescholar@wright.edu.

Amit P. Sheth
University of Georgia

Wil van der Aalst
University of Georgia and Eindhoven University of Technology

Ismailcem B. Arpinar
University of Georgia

This article proposes that an organic workflow-process technology will power the evolution of information system architecture. The authors outline three likely stages of architectural evolution in the context of a networked economy and discuss critical gaps in the current technology with respect to their envisioned future.

Processes Driving the Networked Economy

Speed and distribution will characterize every aspect of most business and organizational undertakings in the next millennium. Organizations will be challenged to bring ideas and concepts to products and services at an ever-increasing pace. Companies distributed over space, time, and capability will have to come

together to deliver products and solutions in the global marketplace. The trends for virtual corporations and e-commerce, along with increased economic global networking, are real and will accelerate. Using the Internet and the Web as the primary communications, interoperability, and integration platform, information systems will play an increasingly critical role in providing a competitive edge for organizations in the networked economy. So far, most of the attention in Information Systems has gone to data. We believe that this attention will increasingly shift to information and knowledge on one hand and processes on the other. The first deals with service and product, the second deals with how to effectively support or render it. This article focuses on the second issue.

An organizational, or business, process is any multistep activity that supports the organization's mission, such as providing a service or manufacturing a product. Today, workflow technology is the most important software technology that supports and automates organizational processes. The research and technologies that make up workflow-process manage-

ment represent tremendous diversity, due in part to the breadth, complexity, and multidisciplinary nature of the problems workflow-process management addresses.¹ Consequently, compared to contemporary technologies such as messaging, database management, and applications servers, it is both harder to provide comprehensive coverage to all the issues facing workflow—a limitation this article admits to—and to build technology that enhances it.

We see process as an organic part of doing business in the future—that is, although processes will chiefly differentiate between the competitive forces in the networked economy, they will be deeply integrated into business itself. Processes will be critical components of almost all types of systems supporting enterprise-level and business-critical activities. Unlike database-management systems (an important category of tools that will continue to be standalone products on which to build applications), workflow technology will not be as significant a market force as a separate category of software tools. For Web sites with further information, please see the related sidebar.

Does workflow technology have a future?

Since the early 1990s, several vendors have offered general-purpose workflow-management systems. The number of workflow products offered peaked at somewhere between 200 and 300 around 1996, and has declined from that point.

Several factors explain this lack of success in today's generation of workflow-management systems. First, many projects introducing workflow-management systems had to deal with legacy problems and the burden of application integration. Workflow-management systems were positioned as a silver bullet that would solve all kinds of problems ranging from technical issues (such as screen scrapping) to organizational issues (such as group self-regulation). As a result, they could not meet the expectations.

Second, the lack of real standards combined with a large volume of vendors has created a scattered landscape where customers are reluctant to invest in workflow products. The numerous workflow-management systems on the market today are based on different paradigms and offer contrasting functionality. Third, most of the production-class workflow-software products offered today are very restrictive and inflexible.

Meanwhile, when the workflow market started to grow in 1992, other market segments have started to co-opt workflow capabilities. Enterprise Resource Planning increasingly supports workflow capabilities. Most leading ERP systems offer a workflow component. Some have developed their own technologies, while others have established partnerships with workflow vendors.

A new market segment that adapted workflow functionality is Enterprise Application Integration. EAI focuses on application interoperability and integration issues within an enterprise. The average Fortune 2000 company relies on 49 enterprise-level applications to run its business and spends 25% to 33% of its IT budget just to get them to talk to one

Web sites for further information

Workflow links (<http://www.workflowsoftware.com>)

LSDIS Web page (<http://lsdis.cs.uga.edu>)

Commerce One's MarketSite.net (<http://www.marketsite.net>)

LIMITrader Securities Inc. (<http://www.limitrader.com>)

SciQuest (<http://www.sciquest.com>)

SigGROUP (<http://www.acm.org/siggroup>)

The Ontology Group (<http://www.ontology.org>)

Trading Edge Inc. Bondlink (<http://www.tradingedge.com>)

VARIA (<http://www.varia.com>)

another. Overall, an estimated 40% of all IT expenditure is devoted to system, application, and data integration. This explains EAI's product appeal, especially in organizations where IT management has a loud voice. Several EAI products currently support limited forms of workflow capabilities through messaging and publish-subscribe mechanisms, but there are signs of future support for more comprehensive workflow-process capabilities. Workflow has become an important differentiator among the products.

Another new software market segment is e-commerce, which is in the midst of explosive growth. Application servers provide support infrastructure for rudimentary workflows. This is poised to become increasingly sophisticated, and given the attention focused on this market segment, new workflow technologies will likely come about as part of new e-commerce products. Given the marketing advantage, even some workflow vendors are in the process of positioning their products in the e-commerce segment.

Looking to the future, we discern two trends. First, vendors are targeting vertical sectors or industry-specific solutions (such as telecommunication, healthcare, distribution, transportation, and central and local government). Domain-specific applications such as call-center packages also reap the fruits of today's workflow technology: Lucent's call-center product uses InConcert to manage call center processes. Second, given the complementary nature of workflow, EAI, and e-commerce, especially in the context of vanishing corporate boundaries in the networked economy and technological advancement in the Internet age, a new breed of products will dynamically cre-

ate and support virtual communities of commerce partners. Current EAI vendors—including Vitria and Active Software—are increasingly adding workflow capabilities to achieve a competitive advantage, while products with roots in workflow increasingly

add EAI capabilities. New breeds of products, such as EAppS from Infocost (which is based on the University of Georgia's Meteor), are taking an integrated approach to combine all three capabilities.

Is it reasonable to say that workflow technology has failed? If we narrowly focus on the workflow-market segment and the predominant vendors from a few years ago, the answer is perhaps yes. However, we see processes as an organic component of any EAI or e-commerce solution. In this sense, workflow-process technology will conduct the emerging networked economy from behind the scenes. Current workflow-process technology can arguably be the basis for building future process support for e-commerce.

Prelude to the networked economy: the telecommunications industry

Arguably the best example of an industry that portends the new networked economy is that of telecommunications services. The telecommunications industry is experiencing a confluence of factors involving globalization, deregulation, and technology breakthroughs. Consequently, not just the technology, but also the businesses are moving at the "speed of Internet." In fact the pace of acquisitions, mergers, and alliances sometimes seems to outstrip the pace of technological changes. The high valuations that Wall Street has afforded to new entrants have fueled entirely new business models, creating a new breed of global corporations that in turn have provided opportunities for applying information

Figure 1. The process portal marketplace architecture.

technology to solve the challenges. Let's briefly review one of the telecommunications industry's main drivers: *convergence* or *next-generation networks*. Generally, this refers to any of the stages of evolution starting from the coexistence and integration of public-switched telephone networks with packet-switched data network, to the development and deployment of a single common packet network for supporting voice, data, video, and other communications services (including what is termed as VoIP—voice over IP). The business driver for convergence is the compelling cost advantage data networks have over switched networks and tremendous growth rate of Internet data traffic (which doubles every four months).

Telecommunications service providers are usually divided based on the services they provide, or the parts of network they primarily own or control. This includes local exchange carriers (LECs) that serve local markets (which are divided into incumbent LECs such as former "baby Bells" and the newer competitive LECs); (b) long distance carriers where competition increased with deregulation; Internet Service Providers including those supporting traditional modem, cable, and xDSL technologies; and wireless service providers. Two of the most critical success factors that this highly competitive industry has identified are customer acquisition and retention, and providing value-added features

and bundled services. Solutions supporting these two compelling needs invariably lead to the need for interorganizational workflow because customers prefer to deal with a single service provider, and most high-value services require integration of what different types of providers have to offer.

Architecture for interorganizational workflows

The Internet offers the ability to transform customer relationships and displace traditional sources of business value because the source of that value is moving from physical products to digital ones. Customers can now choose their own hours of business, access services at any location, and receive attention for their specific needs. In this respect, companies are using the Internet to enter new markets, shrink supply chains, create new value chains, and meet the challenges of global markets.

Using e-commerce to automate inter-business processes across supply chains presents significant challenges. Because custom point-to-point integration between every buyer and supplier is impractical, a possible solution is to transform supply chains into open and interoperable marketplaces.

Some marketplaces host workflow applications, such as those for purchase approvals and accounting. However,

most marketplaces do not have enough facilities to automate the complex business processes conducted between buyers and sellers. In this respect, workflow systems should be exploited to model buying and selling processes.

Terms such as *virtual business process* and *virtual enterprise* clearly demonstrate the relevance of workflows in a networked economy. A virtual business process of a virtual enterprise, also known as interorganizational workflow, goes beyond a single enterprise boundary: it is constructed by combining the services different companies (collectively called a *trading community*) provide. Some of the fundamental issues to address before implementing a virtual enterprise should include how to provide a mechanism whereby companies can advertise their services and how to execute a virtual process that spawns several enterprises without being managed by one physical enterprise.

We define a marketplace as a logically central location in a networked economy where sellers offer their goods or services, and buyers come for convenience and the ability to compare products and prices. Depending on how various stake holders—consumers, intermediaries, and suppliers—interact, and how the capability of managing business processes is realized, we offer three architectures for managing business processes—*process portal*, *process vortex*, and *dynamic trading processes*.

PROCESS PORTAL

A portal is a one-stop shop for products or information. Increasingly, it is also becoming a one-stop shop for services, which are enabled by applications. It takes complex processes involving multiple applications and databases to support or provide such services.

In most current realizations, a portal is responsible for carrying out a majority of activities using the data it has and the transactions it supports. Some of these transactions might involve applications within the organization to which the portal belongs, and a few might even be well-defined transactions or information exchanges with partner organizations (see Figure 1). An information or tradi-

tional e-commerce portals (where selling packaged products is the key activity) primarily involve application servers and databases. A key characteristic of a portal is for it to own or manage much of the data and information it needs to meet customer needs. More advanced portals might use workflow-process support and EAI services to interface with applications, primarily within a single organization. For example, in the telecommunications industry, applications suited for this architecture include consumer-to-business e-commerce applications such as integrated billing and customer self-care.

The same evolution found in other e-commerce activities applies to portals. That is, they first support data (such as Web publishing), then Web hosting and application servers, followed by database and transaction oriented e-commerce. Current portals are expected to evolve to host applications (such as enterprise resource planning systems) so that a company can outsource its application and data management.

Looking to the future, expect to see more advanced forms of portals in the form of *process portals*. A process portal will manage a variety of customizable processes. An individual activity in the process might involve participants or access to information systems at the subscribing company or one of its e-commerce trading partners.

PROCESS VORTEX

The main distinction between a portal marketplace and a vortex marketplace is that interactions between buyers and sellers occur through a third-party marketmaker as opposed to peer-to-peer interactions between buyers and sellers in a portal (a vortex is not a one-stop shop; see Figure 2). Vortex marketplaces generally focus on very specific product lines for specialized markets, and they do not provide access across multiple industries or product groups. Sellers and buyers of these products meet with each other through a particular vortex. However, enterprises might participate in several marketplaces as sellers or buyers.

Sellers advertise their goods and ser-

Figure 2. Process vortex marketplace architecture. Interactions between buyers and sellers occur through a third-party marketmaker.

vices using appropriate tools in the marketplace. The business processes are generally designed to incorporate different trading models (such as auctions), and they are based on structured trading rules. The vortex might also provide process templates for buyers to realize their buying activities. These processes might involve brokering activities such as comparison of goods, prices, and so on. Thus, a vortex provides an organic support for business processes that are usually determined by process builders, usually from vortex sites and supplies. These processes might access the entries of a catalog where the partners within a trading community can post their services and needs.

In the vortex model, the marketmaker converts multiple catalogs from different vendors into a common ontology. Using a unified ontology, buyers can access multiple products from a variety of suppliers and get all the information they need to make a purchase decision. Similar to portal architecture, workflow processes are predominantly predefined. However, these processes can be customized, and they tend to evolve over a period of time.

The *process vortex* architecture becomes relevant in the telecommunications industry when service providers need to support different classes of customers (such as individual residences, small businesses, or large businesses) and

require flexibility to deal with a limited set of partners.

We expect traditional multitier software architecture to be prevalent in implementing a vortex. However, intelligent agents might eventually be used to fully realize a marketplace. Instead of a single buying or selling agent, multiple coordinated agents will dominate the commerce process.

DYNAMIC TRADING PROCESSES

Trading communities focus on very diverse product lines for several markets and provide access across multiple industries and product groups. Many complex and intimate concurrent interactions might occur between peers. In general, participants in a virtual marketplace are a group of semi-autonomous or autonomous organizations that need to cooperate. This requirement is critical because organizations need to preserve autonomy in a competitive business environment—they derive benefits from their unique business and technical capabilities. Therefore, they participate in a virtual business process to gain the benefits of being in a group, but they hide relevant parts of a virtual business process and provide only partial visibility to other partners. Consequently, business relations are more complex, and they are subject to numerous constraints.

Because of their relative independence, business processes can also be

Figure 3. Dynamic trading processes.

highly dynamic in this virtual marketplace architecture (see Figure 3). Unlike portal and vortex architectures, neither business processes nor the set of possible interactions are predefined. Instead, a unique process can be dynamically constructed on a per customer basis. That is, based on a customer's needs and preferences, a virtual process is constructed on the fly to meet a specific demand as depicted in Figure 3. Note that ? and X on the edges within the workflow processes and relationships between companies indicate this dynamic nature of constructing business processes. In general, this process would take each selected company's capabilities and services into consideration. In a more complex scenario, because a customer might wish to deal with a single company to purchase several goods and services at once, that particular company might start a purchasing process. According to negotiations with other companies, missing components of the overall process that meet the customer's needs would be constructed one by one in a way that brings all the pieces of the puzzle together. This shows that a very flexible and adaptable process support mechanism is needed to support dynamic trading processes in a virtual marketplace.

An example of this architecture's application in the telecommunications industry is as follows. One of our visions of future networks includes the facility to allow consumer devices to interact with other devices and humans on the network in an integrated fashion. Here

the device might be able to specify a need for a specific type and quality of network services required, and the network will be able to dynamically compose a customized process to process the request. Additionally, this process might consider the overall value of the process to the customer and the need to maximize the value of the service provided by the business to different customers with different QoS requirements. All the while, the sets of consumers, suppliers, and service demands might be changing, making this a highly dynamic environment.

Key technologies

In the last decade, many software developers and researchers have worked on concepts, methodologies, techniques, and tools to support workflow-process management. Given the networked economy's requirements in the next millennium and the corresponding architectural alternatives discussed earlier, let's discuss some of the key technologies or components of workflow-process management along with the problems they need to solve.

WORKFLOW DESIGN

Many perspectives characterize a workflow design.² The dominant perspectives are process, organization, information, and operation. In the process perspective, *workflow-process definitions* (workflow schemas) are defined to specify which *tasks* need to be executed and in what order. A task is an atomic piece of work.

Workflow-process definitions are instantiated for specific *cases*,³ examples of which include a request for a mortgage loan, an insurance claim, a tax declaration, or an order for a new telecommunications service. Because a case is an instantiation of a process definition, it corresponds to the execution of concrete work according to a specified routing or set of business rules.

In the organization perspective, the *organizational structure* and the *population* are specified. The organizational structure describes relations between *roles* (resource classes based on functional aspects) and *groups* (resource classes based on organizational aspects) and other artifacts clarifying organizational issues (such as responsibility and availability). *Resources*, ranging from humans to devices, form the organizational population and are allocated to roles and groups.

The information perspective deals with *control* and *production data*. Control data are introduced solely for workflow-management purposes (variables introduced for routing). Production data are information objects (such as documents, forms, or tables) whose existence does not depend on workflow management.

The operation perspective describes the elementary operations resources and applications perform. Typically, these operations are used in the process perspective to create, read, or modify control and production data in the information perspective. Most operations are (partially) implemented by applications.

The integration perspective is the link between the other four perspectives (see Figure 4). *Activities* (also called tasks or steps) and workflow-process definitions identified in the process perspective are linked to roles, groups, and resources in the organization perspective, data elements in the information perspective, and operations in the operation perspective. Operations in the operation perspective are linked to data elements in the information perspective, and so on. A *workflow definition* (also called workflow type) is the specification of a workflow that covers all these aspects. Cases are instances of a workflow type and are handled accordingly. A *workflow-management*

Figure 4. The five perspectives in a workflow-management system.

system aims at supporting the five perspectives shown in Figure 4. The build-time service of the workflow-management system allows for the specification of these perspectives. The workflow-management system's enactment service takes care of the actual enactment.

The process perspective is the center of any workflow design. Many languages have been proposed for designing workflow-process definitions. These languages are typically graphical and use building blocks such as OR-split, OR-join, AND-split, and AND-join to model sequential, parallel, conditional, and iterative routing. Figure 5 shows an example of a workflow-process definition modeled with the Meteor builder. The workflow-process definition consists of 13 tasks and specifies the processing of travel requests. Conditional arcs model the causal relations between tasks. The routing conditions depend on control data specified in the information perspective. Meteor supports the explicit modeling of the workflow perspective, but most workflow-management systems do not have such a facility: the information perspectives (and the operation perspective) are modeled implicitly as extensions of the workflow-process definition. The only perspective that is typically modeled in separate diagrams is the organization perspective.

Lack of adequate standards for workflow modeling

In the last decade, many languages have been proposed for the modeling of workflow-process definitions. In fact, in the early 1970s researchers worked on techniques for the modeling of office procedures. Most workflow-management systems use a graphical design language where tasks are connected by arrows and routing elements. Despite the efforts of standardization bodies such as the Work-

Figure 5. Workflow-process definition.

flow Management Coalition (WfMC), there is no real consensus on the representation of workflow processes. Standards such as Interface 1 of the WfMC and PIF (Process Interchange Format) focus on the syntax of the modeling language rather than support specifications or descriptions of process semantics. PIF is an interchange format designed to help automatically exchange process descriptions among a wide variety of process tools such as process modelers, workflow systems, process repositories, and so on. These tools can interoperate by translating between their native process description format and PIF, and vice versa. Other standardization efforts such as Interface 4 of the WfMC, focus on interoperability issues rather than on a uniform design language.

The lack of real standards for workflow management can be explained by comparing the current situation with the early days of database management systems. In the early '70s, most of the pioneers in the field of database management systems were using their own ad hoc concepts. This disorder and lack of consensus resulted in an incomprehensive set of database management systems. However, emerging standards such as the relational data model and the entity-relationship model led to a common formal basis for many database management systems. As a result, their use was boosted.

Until now, such a formal basis is missing from the workflow domain. Many researchers advocate the use of Petri nets as a formal basis for workflow modeling. The Petri net formalism combines a

strong mathematical foundation with an intuitive graphical representation. Several workflow-management systems use a modeling language based on Petri nets (such as COSA, INCOME, BaanERP, and SAP R/3). However, most workflow systems use a vendor-specific diagramming language. These diagramming languages typically abstract from states and do not allow for advanced routing constructs involving a mixture of choice and synchronization. As a result of these and other limitations, it is difficult to translate a workflow-process definition from one system to another. For the organization perspective, there are even fewer consensuses. Some workflow-management systems use a simple role-based mechanism in which there is one work queue for each role. Other systems (such as COSA) provide an advanced scripting language with multiple allocation dimensions that include roles, groups, and authorizations.

Perspectives, languages, and definitions

The technology for designing and exchanging information and operation perspectives is mature and can be used for e-commerce applications and other future workflow applications. However, the organization perspective is underdeveloped compared to the others. In a networked economy, workflow processes will cross organizational boundaries, and these boundaries will become fluid and subject to continuous change. For example, the difference between a customer and an employee is fading. Traditionally,

a customer would be outside of an organizational model, but customers will increasingly track orders over the Web using the same software as employees. Furthermore, in future e-commerce applications, the customer will be able to influence the underlying workflow process by updating requests or supplying additional information. This will require new paradigms to design and manage the organization perspective accurately and effectively.

The limited expressive power of the design languages used by today's workflow-management systems continues to be a source of problems. Most workflow-management systems are not able to model states—they can't model milestones or let the environment select the next step. These systems also have problems dealing with situations involving a mixture of synchronization and choice. In fact, the majority of workflow-management systems use a design language that corresponds to the so-called class of free-choice Petri nets.³ It is well-known that the expressive power of this class is limited compared to standard Petri nets.

Another persistent issue is the reuse of workflow-process definitions. Although workflow processes within different enterprises have common elements, they are typically designed from scratch. Within large companies it is often impossible to specify a workflow-process definition once and replicate it across all parts of the company that need it.

Uniform solutions

Local differences have to be taken into account to prevent the use of one uniform solution to any of these problems. As a result, workflow processes are typically designed from scratch and the wheel is re-invented every day. To avoid this, we can use workflow templates, which is a standard design of a common workflow process. Templates let designers reflect local differences (resulting from specific regulations, organizational structures, and other particularities) and still re-use common parts.

The idea of using templates for workflow processes is not new. Today's ERP

systems offer hundreds of readymade workflow templates (often called business or reference models) that can be used as a starting point for configuring a system. These workflow templates are often based on "best business practices" and reflect the experiences of leading enterprises. Although the set of workflow templates offered by today's workflow-management systems is still limited, we predict the use of templates will increase to avoid starting from scratch every time a new workflow has to be designed. Templates can be shared between different enterprises in the process portal or vortex architectures, thus providing a basis for common understanding and shared business intelligence.

WORKFLOW ANALYSIS

The correctness, effectiveness, and efficiency of the business processes the workflow-management system supports are vital to the organization. A workflow-process definition that contains errors might lead to angry customers, backlog, damage claims, and loss of goodwill. Flaws in a workflow definition's design might also lead to high throughput times, low service levels, and a need for excess capacity. This is why it is important to analyze a workflow-process definition before it is put into production. Basically, there are three types of analyses: *verification*, *validation*, and *performance*.

Verification focuses on *syntactic correctness*; it is independent of the context—it refers to the minimal requirements any workflow should satisfy. For example, there should be no tasks without input conditions. Note that syntactic correctness not only refers to the workflow-process definition's structure but also to dynamic behavior and other perspectives. Examples of syntactic errors in the organizational perspective are roles and groups without any members and cyclic hierarchical relations. Syntactic errors in the integration perspective are pointers to entities in another perspective that does not exist (such as when a task points to a removed role). Potential deadlocks and livelocks are examples of syntactic errors in the process perspective. An important correctness criterion is the so-

called *soundness property* that guarantees proper termination, which means that the predefined final state is reached and that there are no dangling references after termination.³

Validation is concerned with the following question: Is the mapping from the (desired) real-world situation to the workflow design correct? Validation focuses on semantic correctness. To detect semantic errors, knowledge of the application domain is needed. For example, sending a bill before a delivery is confirmed is a semantic error, not a syntactic one. Verification and validation only consider the logical correctness of the workflow processes rather than quantitative measures such as time and costs.

Performance analysis focuses on quantitative measures such as average order lead time, percentage of cases handled within a week, number of cases in progress, and utilization of bottlenecked resources. Today's workflow-management systems give limited support to performance analysis—they provide a rudimentary simulator or a gateway to a simulation tool. Simulation can estimate key performance indicators by experimenting with the specified workflow under an assumed specific environmental behavior. Most workflow-management systems do not give any support for workflow verification. As a result, workflow-process definitions become operational before they are thoroughly checked for correctness. This often results in runtime errors that need to be repaired on the fly at high cost.

Verification of workflow-process definitions

For decades, research groups all over the world have worked on verification techniques. These techniques have checked designs of communication protocols, multiprocessor systems, traffic systems, manufacturing systems, and consumer electronics. State-of-the-art verification techniques let researchers analyze these systems, which typically have millions of states.

The applications mentioned are mainly technical; there are hardly any examples of the use of verification tech-

Figure 6. Verification of a workflow designed with Meteor using Woflan.

niques for analyzing business processes. Business processes are often hidden inside application software, informal business rules, and the minds of the people involved in the process. Moreover, the expertise required to use verification techniques is often missing in the business domain. However, the widespread use of standardized software for business processes (such as workflow-management systems, ERP systems, and BPR tools) is changing this situation. The explicit representation of workflow processes enables the use of verification tools. A typical workflow process consisting of 50 tasks can easily have up to half a million potential states for a single case! Therefore, we need advanced verification techniques. Several groups around the world are working on verification techniques for workflow-process definitions. Most of these groups use techniques based on Petri net theory. One workflow-verification tool, Woflan, can import workflow-process definitions from several workflow-management systems and has analyzed realistic and complex workflows.

Current tools

Both manufacturers and users of workflow-management systems see the need for analysis capabilities. However, few systems have reached a level where verification, validation, and performance analyses are supported in a satisfactory manner. For example, none of the commercially available workflow-management systems offer verification capabilities that go beyond trivial checks such as the absence of an initial task or input condition. In most workflow-management systems, it is possible to model the synchronization (say, AND-join) of two alternative paths (in this case, two paths starting with an OR-split) without any warning at design time. At runtime, such a construct will inevitably result in deadlocks.

Consider, for example the workflow-process definition (modeled with the Meteor builder) shown in Figure 6. Compared to the workflow shown in Figure 5, this figure contains a causal relation between ReserveTickets1 and ReserveTickets2 (task ReserveTickets2 can only be executed if ReserveTickets1 and either CheckApproval2 or Con-

firmApproval2 has been executed with a positive result). This workflow deadlocks if the lower part of the workflow is activated. Moreover, if the upper part is chosen, the workflow will not be able to terminate properly because the trigger from task ReserveTickets1 to ReserveTickets2 is still there, resulting in a dangling reference to a case that does not exist anymore. Figure 6 shows the diagnostics the workflow-verification tool Woflan provided that indicate the error's source. The workflow-process definition made with Meteor was imported by Woflan, which automatically translated the process definition into a Petri net and used various analysis techniques (coverability graph, linear algebraic techniques, and structural methods) to locate the error. Woflan can also download workflow-process definitions from a limited number of other workflow tools. Given the availability of the technology, it is remarkable that today's workflow-management systems offer hardly any verification support. Especially in environments such as the networked economy where changes are frequent and disrupt-

Figure 7: Evolution of workflow-runtime system (enactment service) architectures

tive, the added value of good verification facilities is considerable.

Short-term simulation

Enterprise information systems (powered by workflow-management systems) have knowledge of business processes, the current state of each process, and historical data. This knowledge enables the application of a special form of decision support: *short-term simulation*, which exploits the information that is available. Decisions that affect the business processes in the near future can be evaluated without the need for any additional modeling efforts. The structured storage of information in an enterprise system enables a relatively simple creation of a simulation model. Several workflow-management systems support simulation or provide a gateway to existing simulation tools. For example, ExSpec can simulate COSA workflows, the business processing reengineering tools ARIS and BusinessSpecs can import and export Staffware procedures, and Meta Software's workflow analyzer can interface with Visual WorkFlo and FloWare. These simulation facilities support strategic decision-making.

However, in enterprises where workflow-management systems are used, we also need decision support for management and operational control. The traditional approach toward simulation does not work for this purpose because it focuses on strategic decisions with long-term effects. Research efforts should aim at simulation facilities that

are concerned with the effects of a decision in the near future and exploit all the information available including the actual current state. Such facilities can provide managers with a kind of fast-forward button. By pushing this button, we can see the current situation extrapolated. We can also see the effect of certain decisions (for example, hiring additional employees or renouncing new orders) in the near future. Advanced analysis capabilities for verification, validation, and performance analysis will give managers the tools to react promptly to the ever-increasing pace of change. The applicability of the next generation of workflow-management systems depends on these capabilities.

ENACTMENT

A workflow-enactment service consists of execution-time components that provide workflow processes with an execution environment. It enforces intertask dependencies, schedules tasks, manages workflow data, and ensures a reliable execution environment. Most workflow-management systems and many EAI systems adequately support basic features (mainly scheduling tasks by interfacing them with applications and supporting human participation) and a variety of optional and advanced features such as monitoring, animation, and reporting.

Most workflow systems today employ a three-tier client-server architecture and involve external applications that perform geographically distributed tasks on different platforms. Most existing

workflow-runtime systems are still centralized in the sense that a single workflow engine handles an entire process execution. Unfortunately, this centralized architecture cannot support reliable and consistent process execution with increased failure resiliency, performance, and scalability.

In a fully distributed architecture, the centralized engine is eliminated, and enactment is achieved through geographically dispersed workflow components. In general, the main components of a workflow-enactment architecture are the workflow scheduler and task managers that manage individual tasks on the system's behalf. A workflow scheduler could be either centralized or distributed. Distributed object-based architectures or lightweight agent-based systems provide better infrastructure for dynamic trading processes and for making the process management organic and integral to the systems that will support applications in a networked economy. Figure 7 shows an evolution in the architecture of workflow-enactment services.

Two examples of fully distributed enactment services are the ORBWork and WebWork components of the Meteor system and its commercial version, EAppS. WebWork relies only on the Web for its infrastructure, while ORBWork exploits Web, Java, and CORBA (including some services). Neither have a single entity responsible for scheduling task-manager activation. Instead, the scheduling information is distributed among the individual task

managers. Each task scheduler has the necessary information about its immediate predecessor and successor tasks, thus it is capable of activating the successor task schedulers once the task it controls terminates. Another example of a distributed architecture is *Wide* (Workflow on Intelligent and Distributed database Environments), which is a spin-off of a European Commission project. Wide is a workflow system that provides a distributed environment based on an active database management system to support workflow enactment. It is based on a client-server architecture where the agents working within the system activate clients.

Recently, a lightweight approach was proposed to meet the needs of today's highly dynamic electronic business environment. In this approach, workflow enactment is achieved by a collection of autonomous, problem-solving agents. This approach is explained briefly in the "QoS: Achilles' Heel of workflow management" section.

The next two implementation architectures (distributed cooperating objects or agents and organic process support) are rather new compared to other architectures. In the first architecture, different workflow components can be realized as distributed cooperating objects, and a naming service can be utilized as a way of providing location transparency for these components. For example, all of the ORBWork components—task schedulers, task managers, data objects, and so on—are implemented as CORBA objects, and the ORBWork scheduler is composed of a network of cooperating task-scheduler objects. As we mentioned earlier, processes will become an organic component of any EAI or e-commerce solution in the near future. So, we expect to see a workflow-runtime architecture as an integrated component of future critical-enterprise application servers. Systems from Infocsm, Vitria, and Silknet exemplify such a move.

QoS: Achilles' Heel of workflow management

Anecdotal reports from industry increasingly suggest that systems fail to define,

let alone measure and publish, QoS information. Many products suffer from large process space or virtual memory requirements, while others fail to utilize multithreading or multiple distributed server capabilities in a distributed computing environment. Adoption of workflow management for supporting mission-critical functions in a networked economy will require significant additional progress.

An enactment service should be scalable in terms of carrying large loads. Therefore, partitioning a potentially large load among participating components and guaranteeing minimal communication between these components are essential to provide scalability. The error-handling and recovery framework for a workflow system should also be defined in a scalable manner (for example, by partitioning the recovery mechanism across local hosts).

Transactional aspects such as recovery, atomicity, and isolation to ensure correct and reliable workflow executions are studied in the context of workflow systems. Failures could occur at various stages within the workflow-enactment process's lifetime. Reliability requires that tasks, their associated data, and the workflow system itself be recoverable in the event of failure, and that a well-defined method exists for recovery. In some workflow systems, failure atomicity is ensured through forward and backward recovery. Other traditional recovery techniques such as logging are also successfully used in workflow systems. However, because workflow processes are more complex and fundamentally different from *ACID* (Atomicity, Consistency, Isolation, Durability) transactions, adequate solutions for transactional support in workflow systems are still missing.

The term *transactional* workflow was introduced in 1993 and further elaborated on¹ to recognize transaction relevance in workflows. Transactional workflows support selective use of transactional properties for individual tasks or entire workflows.

Another critical challenge for a workflow-enactment service is its ability to

respond effectively when exceptions occur. In general, an exception can be considered to be any deviation from a process that achieves the process goals completely and efficiently. Workflow exceptions can be categorized as workflow-system-specific exceptions (for example, exceptions due to system failures) and workflow-application-specific exceptions (such as exceptions due to incorrectly performed tasks or when a deadline for a task expires).

Currently, workflow systems provide little support for exception management. These systems typically assume a more or less idealized process. In particular, when exceptions occur, workflow systems mainly rely on the database environment's recovery capability, or users are forced to go behind the workflow system's back, making the system more of a liability than an asset. In some research prototypes, rule-based approaches or artificial intelligence (knowledge-based) techniques detect and resolve workflow exceptions.

The highly dynamic and unpredictable nature of business processes makes agent technology appealing. A workflow-management system architecture can be designed to consist of functionality based reusable components, each of which is realized through different agents. For example, at the lowest level, task agents can be used to invoke different types of tasks. A scheduling agent can be used to support dynamic changes on control flow that emerge from workflow-agent negotiation at runtime.

Organizational modeling and other open issues

Tomorrow's networked economy requires a powerful and reliable execution environment to enact business processes that can span the boundaries of multiple organizations. Significant additional research and serious engineering efforts are needed to improve scalability, exception handling, automatic recovery, and other QoS criteria.

Building all these capabilities from scratch within a workflow system without using any state-of-the-art support tools is not an easy task. In this sense, Iona Technology's OrbixOTM 3.0 with

Articles for further information

In this box, we provide some references for further information for some issues presented in this article.

DOES WORKFLOW TECHNOLOGY HAVE A FUTURE?

D. Sholler, *ENT News and Views*, 9 Dec. 1998, p. 16; <http://www.ent-mag.com>.

PRELUDE TO THE NETWORKED ECONOMY: THE TELECOMMUNICATIONS INDUSTRY

G. Lenahan, "A Practical View of Network Evolution," *Next Generation Networks*, 1998; <http://www.telcordia.com>.

The Secret Sauce of Convergence, J.P. Morgan Securities Inc., New York, 14 Dec. 1998.

Readings in electronic marketplace architectures

G. Alonso et al., "WISE: Business to Business E-Commerce," *IEEE Ninth Int'l Workshop on Res. Issues on Data Engineering*, IEEE Computer Soc. Press, Los Alamitos, Calif., 1999.

S. Arpinar, A. Dogac, and N. Tatbul, "An Open Electronic Marketplace through Agent-based Workflows: MOPPET," *Int'l J.*, to appear late 1999.

Commerce One, *Enabling the Business-to-Business Trading Web Using marketSite 3.0 Open Marketplace Platform*, Commerce One, New York, Mar. 1999.

Comm. of the ACM (Special Issue), Vol. 2, No. 3, Mar. 1999.

G. Dalton, "Bazaar Advantages," *Information Week*, 10 May 1999.

The Ontology Group, *Reference Architecture for iMarkets, Part 1 Overview. Draft 1.1*, Ontology.org, Sep. 1998; <http://www.ontology.org>.

Lack of adequate standards for workflow modeling

W.M.P. van der Aalst, "Three Good Reasons for Using a Petri-net-based Workflow Management System," in *Information and Process Integration in Enterprises: Rethinking Documents*, T. Wakayama et al., eds., Kluwer Academic Publishers, Norwell, Mass. 1998, pp. 161-182.

C.A. Ellis, and G.J. Nutt, "Modelling and Enactment of Workflow Systems," in *Application and Theory of Petri Nets 1993*, M. Ajmone Marsan, ed., Springer-Verlag, Berlin, 1993, pp. 1-16.

PIF Working Group, *PIF: The Process Interchange Format v.1.2*, PIF

Working Group, Dec. 1997, //au: **Web site or publisher location?//**.

WfMC-TC-1016-P, *Interface 1: Process Definition Interchange Process Model*, Workflow Management Coalition, <http://www.aiim.org/wfmc>, Nov. 1998.

M.D. Zisman, *Representation, Specification and Automation of Office Procedures*, PhD thesis, Univ. of Pennsylvania, Wharton School of Business, 1977.

Verification of workflow process definitions

W.M.P. van der Aalst, "Verification of Workflow Nets," *Application and Theory of Petri Nets 1997*, in P. Azema and G. Balbo, eds., Springer-Verlag, Berlin, 1997, pp. 407-426.

N.R. Adam, V. Atluri, and W. Huang, "Modeling and Analysis of Workflows using Petri Nets," *J. Intelligent Information Systems*, Vol. 10, 1998, pp. 131-158.

A.H.M. ter Hofstede, M.E. Orlowska, and J. Rajapakse, "Verification Problems in Conceptual Workflow Specifications," *Data and Knowledge Engineering*, Vol. 24, No. 3, 1998, pp. 239-256.

M. Voorhoeve, "Modeling and Verification of Workflow Nets," in *Workflow Management: Net-based Concepts, Models, Techniques and Tools (WFM'98)*, W.M.P. van der Aalst, G. De Michelis, and C.A. Ellis, eds, //au: **who published this?//**, 1998, pp. 96-108.

QoS: Achilles' Heel of workflow management

I.B. Arpinar et al., "Formalization of Workflows and Correctness Issues in the Presence of Concurrency," *Int'l J. Distributed and Parallel Databases*, Vol. 7, No. 2, April 1999.

D. Barbara et al., "INCAs: Managing Dynamic Workflows in Distributed Environments," *J. Database Man.*, Winter 1996, //Volume? Number? Pages?//.

F. Casati, *Models, Semantics and Formal Methods for the Design of Workflows and Their Exceptions*, PhD thesis, Politecnico di Milano, Italy.

S. Ceri et al., "WIDE: A Distributed Architecture for Workflow Management," *Proc. of RIDE*, //au: **who published this? Where? What pages?//** 1997.

A. Cichocki and M. Rusinkiewicz, "Migrating Workflows," in *Workflow Management Systems and Interoperability*, A. Dogac et al., eds., Springer-Verlag, Berlin, 1998.

C. Hagen and G. Alonso, "Flexible Exception Handling in the OPERA

Java and security support for e-commerce applications seems promising in providing transactional capabilities in a CORBA-based workflow-system implementation. Another promising infrastructure is the one provided by BEA's M3 system.

Organizational modeling in a virtual enterprise and the use of role hierarchies for the enactment of virtual processes are other important research topics. Tomorrow's networked economy might require several organizations to contribute to a single organizational model, and this model might be changed on the fly through a flexible role-domain design tool. Agent-based workflow-management systems still have a long path

ahead before they will effectively address QoS issues.

INTEROPERABILITY

Virtual business processes that operate across enterprises might be implemented using a set of workflow definitions created to support discrete segments of the overall process. To avoid creating islands of automation, different workflow products should talk to each other by exchanging messages that affect process interoperation and integration. Therefore, workflow systems in an organization need to successfully interoperate both internally at the department level and externally with the organizations with which they do business. This can

apply to external parties such as vendors, other businesses, and customers. To achieve wide-scale interoperability between organizations, cooperation between workflow vendors is critical.

The problem of workflow interoperability has many facets, ranging from technological issues about how to integrate different workflow-management systems of different vendors on different platforms to the purely conceptual issues of specifying how the interaction should occur. Unplanned interoperability of different workflows, which are autonomously and separately designed, might cause problems such as deadlocks, starvation, livelocks, or failure to terminate in the desired final state.

Process Support System, *Proc. 18th Int'l Conf. on Distributed Computing Systems*, //au: who published this? Where? What pages?//, 1998.

M. Huhns and M. Singh, "Multiagent Systems in Information-Rich Environments," *Proc. of Second Int'l Workshop on Coop. Info. Agents*, Springer-Verlag, Berlin, 1998.

K.J. Kochut, A.P. Sheth, and J.A. Miller, "Optimizing Workflow: Using a CORBA-based, Fully Distributed Process to Create Scalable, Dynamic Systems," *Component Strategies*, March 1999 //au: is this a journal? What volume, number, month? Pages?//.

N. Jennings et al., "Agent-Based Business Process Management," *Int'l J. Coop. Info. Systems*, Vol. 5, Nos. 2 and 3, 1996.

J. Miller et al., "WebWork: METEOR's Web-Based Workflow Management System," *J. Intelligent Info. Systems*, Vol. 10, No. 2, Mar. 1998.

S. Wheeler et al., "A CORBA Compliant Transactional Workflow System for Internet Applications," *Proc. of the IFIP Int'l Conf. on Distributed System Platforms and Open Distributed Processing*, //au: who published this? Where? What pages?// 1998.

D. Worah and A. Sheth, "Transactions and in Transactional Workflows," in *Advanced Transaction Models and Architectures*, S. Jojodia and L. Kerschberg, eds., Kluwer Academic, Norwell, Mass., 1997.

D. Wodtke et al., "The MENTOR Workbench for Enterprise-wide Workflow Management," *Proc. ACM SIGMOD Int'l Conf. on Management of Data*, //au: who published? Where? Pages?// 1997.

Workflow interoperability specifications

M. Anderson and R. Allen, *Workflow Interoperability: Enabling E-Commerce*, White Paper, <http://www.aiim.org>, April 1999.

F. Casati et al., "Semantic Workflow Interoperability," *Proc. of EDBT Conf.*, //who published, where, what pages?//, 1996.

W. Schulze, C. Bussler, and K. Meyer-Wegener, "Standardizing on Workflow-Management: The OMG Workflow Management Facility," *ACM SIGGROUP Bulletin*, Vol. 19, No. 3, April 1998.

SWAP Working Group, *Requirements for Simple Workflow Access Protocol*, Internet Draft, SWAP Working Group, <http://www.ics.uci.edu/~ietfswap/>, Nov. 1998.

Research on adaptability

W.M.P. van der Aalst et al., "Adaptive Workflow: On the Interplay

Between Flexibility and Support," *Proc. ICEIS*, //au: who published? Where?//, 1999, pp. 353-360.

A. Sheth, "From Contemporary Workflow Process Automation to Adaptive Dynamic Work Activity Coordination and Collaboration," *Proc. Workshop on Workflows in Scientific and Engineering Applications*, //au: publisher? Location? Pages?//, 1997.

F. Casati et al., "Workflow Evolution," *Data and Knowledge Engineering*, Vol. 24, No. 3, 1998, pp. 211-238.

C. S. Ellis, K. Keddara, and G. Rozenberg, "Dynamic Change within Workflow Systems," in *Proc. Conf. on Organizational Computing Systems*, N. Comstock and C. Ellis, eds., ACM, New York, 1995, pp. 10-21.

M. Klein, C. Dellarocas, and A. Bernstein, eds., *Proc. CSCW-98 Workshop Towards Adaptive Workflow Systems*, //au: publisher? Pub. Location? Pages?//, 1998.

M. Reichert and P. Dadam, "ADEPTflex: Supporting Dynamic Changes of Workflow without Loosing Control," *J. Intelligent Information Systems*, Vol. 10, No. 2, 1998, pp. 93-129.

M. Wolf and U. Reimer, eds., *Proc. Int'l Conf. Practical Aspects of Knowledge Management (PAKM'96)*, *Workshop on Adaptive Workflow*, //au: who published, where?//, 1996.

General

A. Dogac et al., eds., "Workflow Management Systems and Interoperability," *NATO ASI Series*, Springer-Verlag, Berlin, 1998.

D. Georgakopoulos, M. Hornick, and A. Sheth, "An Overview of Workflow Management: From Process Modeling to Workflow Automation Infrastructure," *Distributed and Parallel Databases*, Vol. 3, 1995, pp. 119-153.

R. Kalakota and A.B. Whinston, *Frontiers of Electronic Commerce*, Addison-Wesley, Reading, Mass., 1996.

P. Lawrence, ed., *Workflow Handbook 1997, Workflow Management Coalition*, John Wiley & Sons, New York, 1997.

R.M. Lee, "Distributed Electronic Trade Scenarios: Representation, Design, Prototyping," *J. Electronic Commerce*, Vol. 3, No. 1, 1999.

The Ontology Group, *The Need for a Shared Ontology: Business Semantics Underpin XML-Based Architectures*, The Ontology Group, White Paper, <http://www.ontology.org>, 1999.

V. Zwass, "Electronic Commerce: Structures and Issues," *Int'l J. Electronic Commerce*, Vol. 1, No. 1, 1996, pp. 3-23.

Specifications

In general, we can classify interoperability specifications for workflow systems into two categories: specifications for modeling and workflow description, and specifications for runtime interoperability. PIF and the WfMC's Process Definition Interchange (Interface 1) standard fall into the first category, and they provide a standard format for representing workflow specifications. These standard formats exchange workflow specifications between different workflow products. On the other hand, newly emerging products such as the Simple Workflow Access Protocol (SWAP), OMG's jointFlow standards, and WfMC's Interoperability (Interface 4) specification aim to support

exchanging process enactment information or interoperability at runtime. They provide interfaces for operations to support the integration and interactions between workflow systems. Typical operations include operations to create instances of the particular process type and operations to report a workflow instance's status changes. Once a workflow system supports these interfaces, it can talk with other workflow products by exchanging process control and status information through common operations.

The WfMC Interface 4 specification, published in 1998, lets workflow systems from multiple vendors interoperate with one another. It defines the mechanisms that workflow product vendors are

required to implement for one workflow engine to make requests of another. In April 1999, WfMC announced that several vendors had developed the first WfMC Interface 4-compliant products.

SWAP is an Internet-based standard for interoperable workflow products from multiple vendors. The SWAP specification, which is still in draft status, uses HTTP and XML to enable organizations to extend their existing workflow applications out to other organizations on the Internet and Extranets. SWAP is based on jointFlow's object model and the WfMC's workflow architecture.

A group of companies submitted the jointFlow specification in response to OMG's Workflow Management Facility

RFP in 1998. jointFlow addresses the requirements for interoperability between different workflow-system implementations in a CORBA environment.

Open Issues

Although some standardization efforts are under development, interoperability specifications are still far from meeting the interoperability needs of workflow systems, especially those that operate in a virtual business environment. Most runtime interoperability standards make assumptions about the middleware or implementation infrastructures. With multiple organizations participating, the only common ground to be expected is the Internet (HTTP, XML, and possibly Java). Hence the interoperability solutions need to evolve toward multiprotocol and more heterogeneous middleware environments.

In the context of interorganizational workflows, the issue of organizational modeling will become critical. In the future, we expect that workflow specifications will increasingly interact with organization models (including security policies and models, business rules, and so on) that will differ for various participating organizations. The current interoperability specifications do not support organizational aspects in any significant way. Supporting organizational models with respect to the multiple participating organizations involved in dynamic and adaptive workflows will become an increasingly complex issue.

Adaptability

A critical challenge for workflow-management systems is their ability to respond effectively to changes. Changes might range from ad hoc modifications of the process for a single customer to a complete restructuring for the workflow process to improve efficiency. Today's workflow-management systems are ill suited to deal with change; they typically support a more or less idealized version of the preferred process. However, the real runtime process is often much more variable than the process specified at design time. The only way to handle changes is to go behind the system's back. If users

are forced to bypass the workflow-management system quite frequently, the system is more a liability than an asset.

To increase workflow-management system flexibility, it is crucial to know what kinds of changes need to be supported. Basically, there are three kinds of external circumstances that might trigger change: business (motivated by changing markets or individual customer demands), legal (triggered by new legislation), and technological context (introduction of new technology or change in infrastructure). Change can also be triggered by developments inside the system. It is important to distinguish between ad hoc and evolutionary changes.

Ad hoc and evolutionary changes

Ad hoc changes affect only one case or a selected group of cases. The change is the result of an error, a rare event, or a customer's special demands. In general, it is not necessary to change the workflow definition, because the change will most probably not happen in this constellation again. A typical example of an ad hoc change's root is the need to skip a task in case of an emergency. This change is often initiated by some external factor.

Evolutionary changes are of a structural nature: from a certain moment in time, the workflow changes for all new cases that arrive in the system. The change is the result of a new business strategy, reengineering efforts, or a permanent alteration of external conditions. Management typically initiates evolutionary change to improve efficiency or responsiveness, or it is forced by legislature or changing market demands.

Both ad hoc and evolutionary changes are possible at entry time and on the fly. Customizing the process definition for a single case before the processing starts corresponds to an ad hoc change at entry time. If such a customization is also allowed after the processing has started, we name it an on-the-fly change. If evolutionary changes are only possible at entry time, then only the new cases that are started after the change took place have to run according to the updated workflow definition; all other cases run according to the old definition. On-the-

fly evolutionary changes are more difficult to handle because for each running workflow instance, we must decide how to deal with the change.

Problems to solve

Adding flexibility to workflow-management systems is far from trivial. Workflow-management systems such as Ensemble and InConcert support ad hoc workflows, which are defined in an ad hoc fashion by the system's end user before execution. Ad hoc workflow-management systems such as Ensemble and InConcert replicate the process definition for each instance (each case carries its own private workflow-process model). This concept simplifies the handling of change. However, from a management point of view, the replication is less ideal—there is no aggregate management information, and running cases cannot benefit from structural changes without updating every private process definition.

The replication mechanism also degrades system performance. For long-lived workflow processes such as the processing of mortgage loans (which typically have a life cycle of many years), the replication mechanism is unacceptable: a mortgage initiated in 1970 would today have a nearly 30-year-old workflow-process definition. If cases need to be transferred from an existing process definition to a new one, the use of a replication or versioning mechanism will not suffice. The term dynamic change refers to the problem of handling old cases in a new workflow-process definition—how to transfer cases to a new version of the process. We need new concepts and techniques to avoid the anomalies this problem causes.

WE HAVE MANY predictions for workflow technology—they are as much based on our view of various business or industry trends as they are on technology and research trends. The industries where we can first validate our observations are those that are seeing rapid changes due to deregulation, rapid technology changes, and other reasons. The best examples include telecommunica-

tions, financial services (including insurance), and energy services. The laggards include healthcare and heavy industries where regulation and legal concerns combined with the slow pace of technological changes and lack of new investments have limited retarded growth or change. We will see the following:

- Workflow-process management functions and technology will be absorbed by other technologies. Instead of standalone workflow-management systems on which workflow applications are built, workflow capability will be built in critical enterprise application systems such as ERP and supply-chain management. This capability will be supplied as an integral part of a future generation of EAI and other middleware services. We believe that process management will be necessary for products in this area in the near future, just as data access and transaction management are critical to most middleware products today.
- Adaptability will become a key requirement. This increasingly dynamic environment will lead to various kinds of change ranging from ad hoc modifications for an individual customer to evolutionary changes as a result of BPR efforts. Today's workflow-management systems have problems dealing with change and are too rigid to handle the dynamics of the networked economy. Therefore, new concepts will need to enable adaptability without losing control.
- Capturing processes will be one of the main outputs of business intelligence and knowledge management activity. There will be a shift from data-centric to process-centric knowledge. Understanding and managing the dynamics of organized activity will be of prime importance in the new millennium.
- Outsourcing of workflow management will become an attractive option. Currently, corporations already outsource some of their data-related functions as well as their Web sites. Recently, new companies are targeting application

Figure 8. Outsourcing of data, web, applications, and processes.

outsourcing. For example, rather than managing its own ERP applications, an organization might rent an application from (and store the data at) an application service provider and have access to it using a virtual private network or some other means. Organizations desire to concentrate on core competencies will lead to outsourcing process management, especially to support interorganizational processes. Figure 8 illustrates how the outsourcing of processes will have a considerable impact on the way organizations operate.

References

1. A. Sheth et al., "Report from the NSF Workshop on Workflow and Process Automation in Information Systems," *SIGMOD Record*, Vol. 25, No. 4, 1996.
2. S. Jablonski, and C. Bussler, *Workflow Management: Modeling Concepts, Architecture, and Implementation*, Int'l Thomson Computer Press, Stamford, Conn., 1996.
3. W.M.P. van der Aalst, "The Application of Petri Nets to Workflow Management," *J. Circuits, Systems and Computers*, Vol. 8, No. 1, 1998, pp. 21-66.

Amit Sheth directs the Large-Scale Distributed Information Systems lab at the University of Georgia, and is a full professor in the Department of Computer Science. He also founded and serves as president of Infocsm Inc. (<http://www.infocsm.com>). His technical interests include enterprise integration,

semantics in global information systems, digital libraries, e-commerce, and digital media. He received his BE in electrical and electronics engineering from the Birla Institute of Technology and Science, Pilani, India, and his MS and PhD in computer and information technology from Ohio State University. He is a member of the IEEE and ACM. Contact him at the LSDIS Lab, Dept. of Computer Science, Univ. of Georgia, 415 Graduate Studies Research Ctr., Athens, GA 30602-7407; amit@cs.uga.edu; <http://lsdis.cs.uga.edu>.

Wil van der Aalst is an associate professor in the Department of Mathematics and Computing Science at Eindhoven University of Technology. He leads a research group working on software engineering, Petri nets, and workflow management. He also works for the Dutch consulting firm Bakkenist, where he works on workflow-related projects. His research interests include simulation and verification of business processes, Petri net theory, information systems, and logistics. Contact him at the Dept. of Mathematics and Computing Science, Eindhoven Univ. of Technology, P.O. Box 513, NL-5600 MB, Eindhoven, The Netherlands; wsinwa@win.tue.nl.

Ismailcem B. Arpinar is a post-doctoral research Fellow at the Large-Scale Distributed Information Systems lab at the University of Georgia. His primary research work and interests are in the area of database systems, particularly workflow systems, transaction processing, electronic commerce, process repositories, and distributed databases. He received his BS, MS, and PhD degrees in computer engineering from the Middle East Technical University in Ankara, Turkey. Contact him at the LSDIS Lab, Dept. of Computer Science, Univ. of Georgia, 415 Graduate Studies Research Ctr., Athens, GA 30602-7407; budak@ainge.cs.uga.edu; <http://orion.cs.uga.edu:5080/~budak>.