Emerging technologies that support a software process life cycle

by G. T. Heineman
J. E. Botsford
G. Caldiera
G. E. Kaiser
M. I. Kellner
N. H. Madhavji

The goal of developing quality software can be achieved by focusing on the improvement of both product quality and process quality. While the traditional focus has been on product quality, there is an increased awareness of the benefits of improving the quality of the processes used to develop and support those products. These processes are key elements in understanding and improving the practice of software engineering. In this paper, existing objectives for the development and application of models of software processes are restated, and current research sponsored by the IBM Centre for Advanced Studies (CAS) is discussed as it applies to furthering each of the objectives. A framework is also presented that relates the research work to the various sectors of a software process life cycle. The on-going research involves four universities, CAS, and collaboration with IBM Toronto Laboratory developers.

The primary concern of the software engineering community is the development and support of quality software. The two basic approaches toward this goal (see Reference 1, for example) are improving product quality and improving process quality. Product quality focuses on end deliverables and is associated with such concepts as rate of fault occurrence, mean time to failure, and other measurable quantities. Known methods exist to improve product quality, such as code reviews. Process quality focuses on the processes used to produce end deliverables, and is concerned with such topics as accuracy (the degree to which the product produced by the processes matches the intended result) and fitness (the degree to which the people involved in the processes can faithfully follow specified actions). If we consider code reviews within a particular process, we see that the process is greatly affected by such things as: when the code reviews should be performed, how they will be done, who should participate, and how the results should be applied.

While the traditional focus has been on product quality, there is an increased awareness of the benefits of improving the quality of processes. The International Organization for Standardization (ISO) has created an ISO 9000-3 standard[2] specifically for the software industry. This standard defines a life-cycle quality system governing the development and maintenance of software. The Malcolm Baldrige National Quality Award[3] has "Management of Process Quality" as one of its seven assessment categories. Finally, the Capa-

bility Maturity Model (CMM),[4-6] produced by the Software Engineering Institute (SEI) at Carnegie Mellon University, defines increasing maturity levels of an organization based on the quality of the processes employed. This trend of heightened process awareness can be viewed as the maturing of software engineering as a discipline.

In recent years, the software engineering community has come to recognize the importance of the processes that are employed to develop, support, and maintain software. The following objectives[7] are cited as the motivation for developing and applying models of software processes:

- Facilitate human understanding and communication
- Support process improvement
- Automate process guidance
- Automate process execution support
- Support process management

Each objective has an associated list of subgoals called a goal cluster. As these objectives are restated throughout the paper, we discuss current research at the IBM Centre for Advanced Studies (CAS) against these goal clusters. These goals are not an exhaustive list, but they do provide discrete markers in software development processes by which one can measure the progress and impact of current studies.

Before discussing the process study, we define some terms that might not be clear to readers unfamiliar with software processes; most of these are taken from Reference 8 where Feiler and Humphrey define a set of essential terms and concepts about software processes.

- A *process* is a set of partially ordered steps intended to reach a goal. The goals for software development processes include the production, or enhancement, of quality software products. Other software processes include maintenance.
- A *process model* abstracts and captures those aspects of a process relevant to the modeling formalism used. Any abstraction of a process can, in fact, be a process model, but process models are most useful when they can be analyzed, simulated, and validated. They can also be used to aid process understanding.
- *Agents* are the entities that execute a process model by carrying out individual process steps.
- *Automation* is the use of machine agents to perform individual steps.
- An environment *enacts* a process by providing automation and guidance to the agents carrying out the process while enforcing any process constraints. An enacted process is an active process.
- An *enactable process* is instantiated from a process model and contains all the information necessary for an environment to enact the process.
- *Organizational structure* is the configuration of people and other resources that perform activities within an environment and the relationships between them.
- *Process management* involves all activities that plan, control, and manage processes.

This paper describes the joint work of five groups, a consortium of faculty and students from three universities (Columbia University, the University of Maryland, and McGill University), the Software Engineering Institute (SEI) at Carnegie Mellon University, and CAS. Although the work was originally named the Process Reuse Study (PRS), the application of the techniques and ideas produce a holistic vision of software processes by defining a life cycle and a method for designing, improving, and performing software process models. Product improvements can be achieved by constructing models from the actual processes, improving the models, and then implementing the improved processes in practice.

We first describe the PRS framework and in separate sections discuss how PRS addresses each of the five objectives. We conclude with a discussion of the contributions of PRS to-date and plans for future work.

Process study

Effective software processes are one of the most significant assets of a large software development organization; unfortunately, they are often undervalued. The basic premise of PRS is that the quality of a software product is largely determined by the quality of the processes used to develop and maintain it.[9] Describing, studying, and improving software processes can improve the quality of the software and the way the software is developed. Software processes must be as meticulously maintained as the software that they have produced, since they might need to change over time.

Figure 1 The process cycle

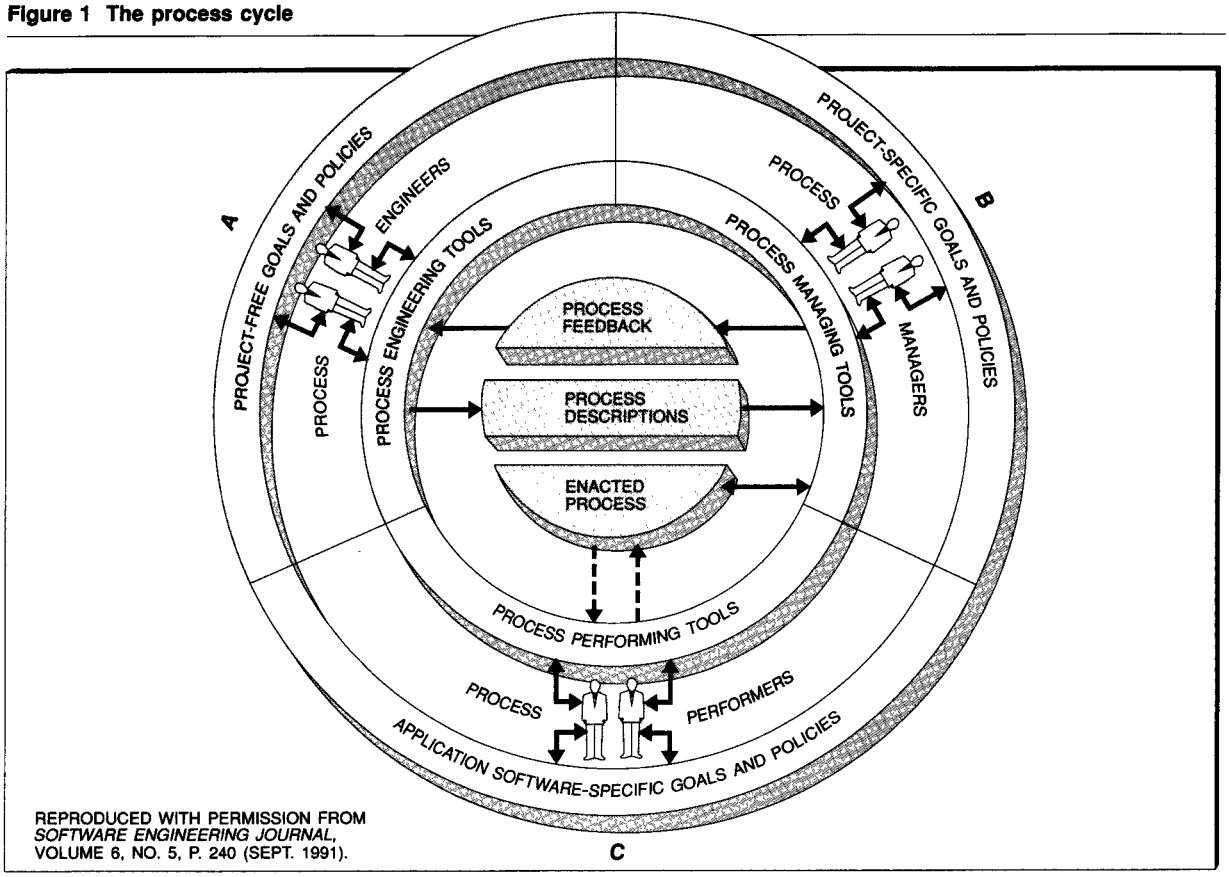

What is needed is nothing less than a life-cycle view for software processes.

A life-cycle view has been developed[10] containing four parts: description and definition, customization and instantiation, enactment, and improvement. From these steps, a process cycle is created that defines the key roles played by humans, categories of tools used, goals and policies governing the process, and interrelationships and feedback among the different roles. The process cycle shown in Figure 1 defines the scope of all process steps necessary for the development and evolution of software processes. To realize this process cycle, PRS has developed a framework, illustrated in Figure 2, that is an implementation of the process cycle. We next describe in more detail the process cycle and the framework.

Process cycle. The process cycle is divided into three sectors (as seen in Figure 1) that represent the engineering, managing, and performing of software processes. In sector A, process engineers design, construct, and improve *generic process models* (or *models* when the context is clear). These models are generic in the sense that they have not yet been tailored for a particular software project. The models are customized by process managers in sector B and are instantiated for use; process performers in sector C carry out the processes.

Two essential concepts to the process cycle are the tools used by each sector (see Table 1) and the feedback among sectors. The success of the process cycle depends heavily upon the success of these tools. It is expected that no single tool can be used in all sectors. Some tools are very adept at modeling and simulating processes, but have no means of enactment, while other tools created to enact processes might be weak in modeling capabilities. The ability to integrate a heteroge-

Figure 2 The process study (PRS) framework

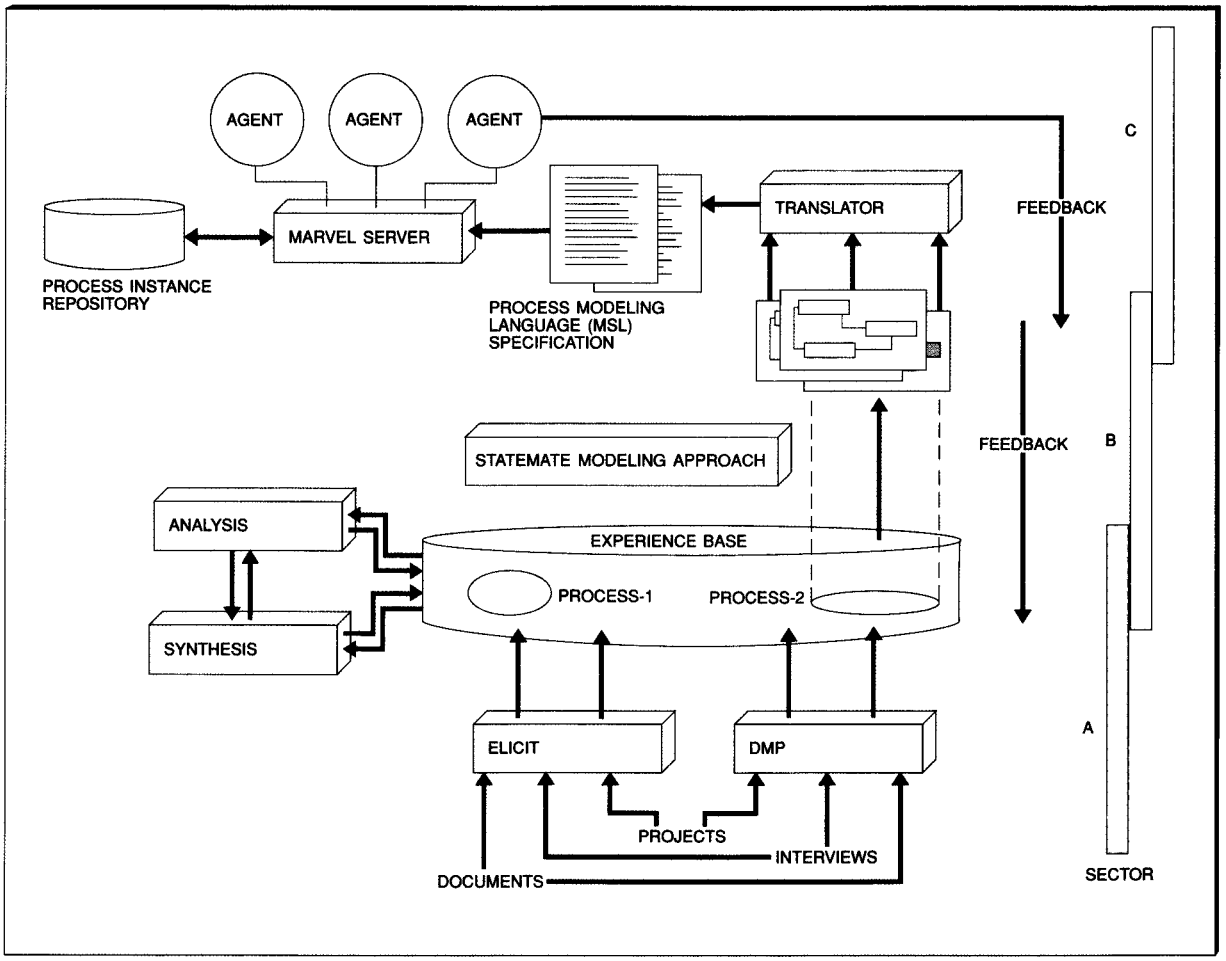

Table 1 Process cycle tools

Sector	Tools used	Example
A	*Tools to create, and operate on, fragments of processes within simulated environments*	Elicit, DMP
B	*Tools to operate on specific descriptions within simulated and actual environments*	Statemate, MARVEL
C	*Tools that operate on the software parts being constructed in actual environments*	CC, VI

neous set of approaches and tools is essential to an effective implementation of the process cycle.

The second important aspect of the process cycle is the feedback among sectors that can be used in each sector to improve processes. Feedback from process performers in sector C might reflect the smoothness with which assigned tasks are carried out, bottlenecks in the development process, negative effects caused by certain timing constraints, and the need for additional process steps or removal of superfluous ones. Both qualitative and quantitative data about processes are useful. Process managers can use feedback in several important ways. First, project-specific processes can be improved through modifications in response to the feedback. Second, generalizations and improvements in models can be suggested to the process engineers in sector A. Third, an iterative approach (see a later section, "OPT approach," for example) can be used to identify the appropriate changes to be made to organizational structures and development processes to satisfy project goals. This cycle of building models, tailoring them, and improving them through feedback is continuous.[11]

PRS framework. The PRS framework, illustrated in Figure 2, is our proposed implementation of the process cycle. The approaches (on the left side of the figure) are matched with an appropriate sector of the process cycle on the right. For example, the Elicit approach occurs within Sector A. We next outline the PRS approach to software processes by describing each of the entities of the framework.

Process engineers first model a process by observing existing processes within an organization. The Elicit and DMP (Descriptive Modeling Process) projects (discussed in a later section) address the aspects of understanding an organizational structure and its processes by extracting the appropriate information from documents and interviews. As models are created and tailored, they are stored in an *experience base*. This repository contains the models and their histories—all modifications to models, lessons learned, and decisions and alternatives.

The experience base is derived from the component factory concept proposed by Basili et al.[12] This repository has three levels of representation that contain descriptions of the agents in the organization, the activities they perform, and the processes used. The three-tiered structure of the repository suggests that the modeling formalism used should have some mechanism for supporting multiple levels of abstraction.

The SEI-developed Statemate***[3]-based modeling approach[13–15] models software processes through formal visual languages based on *hi-graphs*,[16] a form of diagram using boxes, directed lines, and hierarchical nesting. When a process is modeled, three forms of charts are produced, each providing a separate yet interrelated view of the process. *Activity charts* capture the functional perspective of the process and focus on the activities being performed while hiding the actual details of how they are carried out. *State charts* provide the behavioral perspective and focus on the timing and ordering of the individual process steps. Finally, *module charts* describe the organizational units involved and the physical communication channels used to transfer information between the activities. These three charts provide an integrated vision for process models and correspond directly to the three levels of the experience base.

Statemate[16] provides the analytic and simulation capabilities required by process engineers. With this tool, engineers can step through a modeled process, run batch simulations, generate events, and observe the reactions of models. Recent work[17] shows how Statemate can be used to provide quantitative estimates of process time and resource allocation. Process managers can correlate these estimates with statistics derived from the actual performance of processes. Because Statemate offers little help in performing processes, PRS incorporates the MARVEL process-centered environment.[18]

In MARVEL , a process model is defined by a process modeling language (MSL), with each process step encapsulated by a rule. Each rule is composed of a condition, an optional activity, and a set of effects. The condition is a prerequisite that determines when the process step can be carried out. The effects are the immediate implications of the process step. Each process model contains an object-oriented data schema that defines the process state and the organization of the particular product data. MARVEL provides process automation through backward chains that attempt to satisfy the prerequisites of a process step, and forward chains that carry out all implications of a process step. Recent work has shown how Statemate charts can be automatically translated

Table 2 Goals for human understanding and communication

	Goal
1	*Represent processes in a form understandable by humans*
2	*Enable communication about and agreement on software processes*
3	*Formalize processes so that people can work together more effectively*
4	*Provide sufficient information to allow an individual or team to perform the intended processes*
5	*Form a basis for training the intended processes*

into MARVEL rules,[19] thus producing a multiuser environment that can enact a given process model. Sector C tools can be incorporated into the MARVEL environment, and measurements can be taken as processes are performed, providing useful feedback to process managers.

In the remainder of the paper, we discuss PRS research within the context of our original objectives. At the end of each section we assess how well PRS research addresses the goal clusters for the objectives.

Human understanding and communication

The first step toward our goal of improving software processes is understanding them. Table 2 presents five motivating goals for employing models to understand software processes. Industrial software processes often require more than a year to carry out a development or support cycle,[20] making it inconvenient to acquire understanding through direct observation. Because software processes are extremely complex, managers and participants frequently lack a broad sense of intellectual control over them. Even simple processes quickly become complicated because there can be many processes performed simultaneously by different people, so that often no one person understands the entire set of processes. Consequently, gaining understanding is a valuable goal in its own right.

The goal cluster in Table 2 focuses on effectively communicating the description of a process to others, such as workers, managers, and customers.[21] Process developers and researchers must always remember that their goal is to help people understand processes. Models and definitions are simply vehicles for this purpose and are useless if they themselves cannot be understood. A complicating factor in facilitating understanding is that so many diverse groups access models, including software process engineers, project managers, software engineers, system engineers, software executives, and customers. These different groups may place widely differing demands on a modeling approach because of their different information needs and expertise. Visual approaches, good use of abstraction, and multiple perspectives offer promising techniques for coping with these challenges. We next describe the two PRS techniques used to construct models from existing processes: Elicit and DMP.

Elicit. It is our belief that valid models are most easily derived from existing processes. Reference 22 describes a methodical approach, named Elicit, that elicits models from active processes. Elicit is actually a meta-process for eliciting software process models with the following steps:

1. Understand an organizational environment
2. Define elicitation objectives
3. Plan the elicitation strategy
4. Extract process information
5. Synthesize and translate models
6. Perform the analysis

Elicit progresses from an implicit state (undescribed, enacted processes) to an explicit state (described, analyzed models). The derived models are collected and made part of the experience base for the organization, ready to be used for further understanding and improvement. Eliciting software process models from enacted processes involves all the activities that are necessary in extracting process information from various sources and synthesizing it into appropriate models. This elicitation is similar to bottom-up process re-engineering; they both identify the work that an organization should perform and define efficient and effective processes for that organization. Elicit is an evolutionary method that produces feedback that might lead to the iterative repetition of previous steps.

Elicit has been applied and documented in several laboratory and industrial environments. The basic concept of Elicit is discussed in detail in Reference 22. At the IBM Toronto Software Solutions Laboratory, Elicit has been applied to the lab-wide recommended requirements process.[23] Based on this experiment, Elicit was then analyzed and improved.[24] Formalizing, applying, measuring, analyzing, and improving Elicit was the subject of a master's thesis research project[25] and studies in a laboratory environment to elicit a literature assimilation process.[26]

In an early application of the process cycle, a semiformal process notation was designed[27] to be used as a vehicle for understanding and communicating process information. Process notations were developed in each sector of the process cycle using a specification-implementation paradigm. Further work on the infrastructure needed to support the communication of processes is outlined in Reference 28. These human understanding and communication goals were in fact at the center of an earlier joint IBM–McGill project on software processes[29] that led to the recommendation of a process modeling method and of a supporting tool. We next describe in detail the six steps that Elicit encompasses.

Understand an organizational environment. The first step of the meta-process is to understand the environment from which models need to be elicited. This step is of fundamental importance since it defines realistic objectives upon which the rest of the elicitation process is based. The following partial list illustrates some of the questions that can be answered by this step:

- Projects of concern: Is the process to be examined project-specific, multiproject, or organization-wide?
- Team structure: Are the teams organized hierarchically or democratically? Are they autonomous or authoritarian?
- Degree of process awareness: How strong is process awareness among the various groups and individuals in their respective roles? What purposes would models serve? Are the expectations from described processes unreasonably high? How soon are the process descriptions required? What quality of process descriptions is required?

The input to this step is organization knowledge; its outputs are documented contexts considered useful for the elicitation tasks that we next describe.

Define objectives. In this step, the specific objectives of the elicitation task are defined. These objectives are used to measure the quality of the elicited models. Without such objectives, it would be difficult to control the elicitation project. Some objectives to be considered are software process domain, granularity, consistency, completeness, cost, urgency, and resource constraints.

The software process domain is especially important because it provides a setting within which particular elicitation objectives achieve the greatest impact. The appropriate domain should be chosen based on many factors. Some of the areas addressed by this step are:

- Known process problems: Which parts of the process are ill-defined or not well understood? Which parts of the software process contribute to software quality problems? These weak process areas are revealed by process assessments.
- Expected process changes: For example, if a new set of tools is to be added to the design process, this process may need to change to accommodate the new tools.

Granularity and consistency determine the amount of effort to be undertaken by setting the level of details of the elicited models. The more fine-grained a process model needs to be, the

more effort will be required to produce it. Completeness also has an impact on the cost, since it determines the scope of the elicitation. Urgency and resource constraints determine how soon the elicited model needs to be produced, and whether a given set of human resources will be available to complete the desired elicitation.

Plan elicitation strategy. This step builds an elicitation plan and allocates appropriate resources so that when the plan is executed, the desired quality of the elicited process model is developed within budget and on time. This activity involves identifying the scope of the elicitation task in detail, scheduling interviews and document-understanding sessions, identifying the deliverables, selecting elicitation tools, and allocating computing and human resources. This planning step can be as complex as the management of software projects.

The elicitation strategy is affected by both the objectives and the contexts from the previous two steps. For example, the objectives may determine a budget and time schedule that limits the number (and type) of people to be interviewed and documents to be examined. In addition, the documented contexts contain knowledge about particular documents and specific people that can help plan and schedule elicitation activities. The role of the persons planning the elicitation strategy is important and should be carried out by highly skilled and influential process experts. An ineffective plan is likely to result in low-quality descriptive models and wasted time and effort. Also, if such low-quality models are used in a software project for training practitioners, or as a basis for process improvement, then a great deal of damage can occur.

Since Elicit is iterative, the elicitation plan might require either the objectives to be redefined or the contexts to be described more clearly. This step produces a plan to be used for elicitation of models from the identified processes.

Extract software process information. The plan created by the previous step is carried out, including reading process documentation, holding interviews, analyzing question responses, analyzing the elicited model, and demonstrating the elicited process descriptions to management. A tool that supports the meta-process (the Elicit tool) prompts the user at each step for process information that it represents as process models. In essence, Elicit takes care of the mechanical aspects of eliciting, leaving the user with the creative tasks of providing appropriate information. This can save considerable elicitation time and effort and reduces the chance of missing information or having erroneous information.

Once the information has been elicited, a model is created from the structured process descriptions. At this point, static semantic checks can validate the elicited model against the objectives defined previously. In iterative fashion, the previous steps might need to be re-executed to further refine the contexts, objectives, and elicitation plans. This step produces a descriptive process model.

Synthesize and translate. Once the elicited process model has been statically reviewed, its dynamic behavior needs to be examined for correctness. This could be done as part of the previous step if the eliciting tool has behavioral analysis capabilities. Instead of building dynamic analysis and simulation capabilities into our elicitation tool, however, we simply translate the elicited model into a representation suitable for dynamic analysis by a commercial tool. Given this approach, the two key issues for translation are: (1) the compatibility of the modeling formalisms supported by the elicitation tool and the simulator; and (2) the method of translation (that is, hand or automatic translation) of the elicited process model.

The first issue requires the representation scheme supported by the simulator to be at least equivalent to, or a superset of, the one supported by the elicitation tool, otherwise information will be lost during translation. The second issue implies that either an automatic translator is constructed, or careful hand-translation of process components is carried out, followed by reviews. To avoid the complications of translations, Elicit allows the process engineers to hand-translate the elicited information into Statemate charts that can then be analyzed using Statemate.

Analyze. The final step of Elicit is the analysis of the behavioral aspects of the translated process model. If necessary, the model is corrected to reflect the dynamics of the enacted software process. For this purpose, there is a need for appropriate process analysis and simulation tools. The

process models should be checked for deadlocks or race conditions, starvation among subprocesses, reachability problems, idle time and delays in the process, and behavioral ambiguities.[14,30]

The process engineer should also observe the degree of parallelism in process subcomponents, obtain animated traces of simulated paths, experiment with "what-if" scenarios, and calculate quantitative predictions. Researchers have shown how software process models can be analyzed to make prescriptive improvements.[17,31–33] Here, the objective is to make descriptive improvements to the elicited model that reflect the behavior of the enacted process. Prescriptive and descriptive improvements can be made at the same time.

In iterative fashion, all previous Elicit steps might need to be re-executed if flaws are found in the process model. Once this step has terminated, Elicit is complete.

Package. The fulfillment of Elicit occurs during its packaging phase. Once the information has been elicited, the constructed model must be stored in a meaningful manner. Packaged information related to the elicitation task can be reused in other tasks (such as new elicitation tasks or process improvement efforts) and can be helpful for educating process users. Packaging is particularly worthwhile if there are on-going or numerous elicitation (or process improvement) efforts in the organization. There are several key perspectives to consider when packaging the elicited information: objects of interest to package, users of the packaged information, and reusability of the packaged information.

The objects of interest can be categorized into either products of Elicit or experience gained in eliciting software processes. Clearly, we want to store the models created by Elicit; we also store the experience gained, including the contexts, objectives, and elicitation plans constructed by Elicit. This ability to store all pertinent information increases the integrity and cohesion of all elicited models. The users of the packaged information are varied; they include process educators, process improvers, managers, and elicitation-task performers. For example, process educators can use elicited models for software process-understanding tutorials, while process improvers can use the elicited process model and organizational environment as a basis for process improvements. Project managers can obtain reports to create cost benefit analyses while elicitation-task performers can use the packaged strategies in other elicitation tasks. The number of users and the ease with which they access the packaged information increases its reusability. For example, organizational environment details would be extremely useful in eliciting other software process models within the same software project.

The process modeling formalism needs to be flexible enough to satisfy the different needs of the users who will be accessing the packaged information. Two features we have found useful in modeling formalisms are multiple levels of abstraction and multiple perspectives. The users should be able to start understanding the process models from the top working down, from the bottom working up, or even from the middle working either way.[14,30] Furthermore, representing multiple perspectives on a process can be quite beneficial; these may include functional, behavioral, organizational, and informational perspectives.[7] Capabilities, such as levels of abstraction, precision, and multiple perspectives, offer distinct advantages over representations based upon narrative text or a single type of diagram. PRS envisions that all elicited models will be stored in an experience base, whose structure we next describe.

Experience base. The experience base is derived from the component factory[12] that attempts to increase the quality and productivity of software by targeting three goals: improving the effectiveness of the software processes, reducing the amount of rework, and reusing life-cycle products. The production of software using reusable components is a significant step forward for all three of these goals, but there are still problems in achieving higher levels of reuse. The current inability to package experience in a readily available way prevents the transfer of experience from one project to another. Another problem is the difficulty in recognizing experience that is appropriate for reuse. Finally, reuse needs to be an integral part of software development processes before it can be truly effective. The experience base addresses these concerns by applying the concept of the component factory to the domain of software processes.

Like the component factory, the experience base has three levels of abstraction representing the different aspects of a process. The highest and most abstract level, *Reference*, describes the agents in the organization. The *Conceptual* level represents the interface of the agents and the flows of data and control among them. The lowest level, *Implementation*, defines the actual implementation, both technical and organizational, of the agents and of their connections specified at the Conceptual level. These three levels correspond exactly to the three perspectives modeled by the Statemate-modeling approach described in the earlier section "PRS framework."

Development of models increases the understanding of processes.

The experience base is also similar to the Process Asset Library (PAL)[34] prototype developed in 1992 by SEI in collaboration with the STARS (Software Technology for Adaptable, Reliable Systems) program of the U.S. Department of Defense. The PAL prototype includes nine processes, each represented by a model, and seven represented by a process guide as well. These nine assets were primarily based on pre-existing process documentation that included an IEEE (the Institute of Electrical and Electronics Engineers, Inc.) standard, process guidebooks, and a journal article. The PAL prototype contains over twelve hundred pages of process documentation (models and guides) and was the first publicly available collection of industrial-strength models.

Descriptive modeling process (DMP). Since 1987, the group at SEI has developed more than a dozen models of industrial-scale software processes. Several of these are descriptive models of processes as they are actually performed in a specific real organization.[13,14,35] This experience has been coalesced into a Descriptive Modeling Process (DMP) that has been taught in detail to over 200 software professionals for use in their organizations. This three-day workshop offered by SEI will be documented in a forthcoming paper. DMP produces (semi-) formal models of current processes as practiced in a specific organization. Considerable effort and care are necessary to accurately reflect organizational practice, as opposed to the process documentation that records what process is desired. Developing such models involves extensive interviews with process performers to elicit the information and verify the models. The primary inputs to DMP are existing process documentation and knowledge of the as-practiced processes from those who perform and manage them. DMP creates a descriptive model of the as-practiced processes in a variety of representation languages, including Statemate, ETVX (Entry-Task-Validation-eXit), and IDEF0. Most descriptive process models are developed in an iterative, top-down fashion. As an example, one comprehensive model developed at SEI involved interviews of approximately 25 individuals and 5 iterative rounds of construction activity.[35]

Both DMP and Elicit create process models that clearly describe the processes and can be used to aid the performers of the processes (Goal 4 in Table 2) to train new members in the processes (Goal 5 in Table 2).

Achievements in understanding. Formalizing existing processes increases the understanding of the processes. We have found that modelers always increased their understanding of processes through the experience of developing models. More importantly, when process participants, managers, and other personnel viewed a model, they almost invariably reported increased understanding (Goal 1 in Table 2). The impact that models have on process understanding shows their communicative power. Personnel ranging from software engineers to senior managers have found these models to be understandable and useful. The models have allowed them to clearly visualize how the various process components are interrelated. Their comments, questions, and insights arising during presentations of these models demonstrate that models are valuable communication vehicles. Process participants have gained a deeper understanding of portions of the processes in which they do not directly participate, and management personnel have gained substantial understanding as well.[14,23,35]

Table 3 Goals for process improvement

	Goal
1	*Identify all the necessary components of high-yield software development or maintenance processes*
2	Reuse well-defined and effective software processes on future projects
3	*Compare alternative software processes*
4	*Estimate the impacts of potential changes to a software process without first putting them into actual practice*
5	*Assist in the selection and incorporation of technology (for example, tools) into a process*
6	*Facilitate organizational learning regarding effective software processes*
7	*Support managed evolution of a process*

The following analysis of the understanding and communication benefits taken from a specific SEI modeling effort[14] shows how Goals 2 and 3 in Table 2 are addressed. In this study, the focus was on the process of identifying and making changes to Technical Orders[36] (TOs) in response to changes in the software—in particular, changes to the Operational Flight Program for the F-16A/B multirole fighter. A typical new release of the Operational Flight Program required changes to 3000 pages of TOs spread over as many as 100 separate documents.

> Our efforts at modeling the TO Modification Process led to a number of very important results. The first of these is a notable increase in understanding of the process by those involved in executing and managing it. . . .
>
> Several different organizational subunits are involved in various stages of the TO Modification Process. Not surprisingly, the view of most of these subunits was somewhat parochial, focusing on their specific subtask with little awareness of overall implications. . . .
>
> In the course of our interviews, we gained an increased appreciation of the overall goals of the process, as well as a recognition of the effect of regulations and standards on the process. Communication of this information to relevant personnel can result in significant benefits. A common understanding of the overall process and the role each group's work plays in its successful completion, may lead to increased goal congruence among those involved. . . .
>
> We also found that graphical representations of the process were far more effective communication vehicles than narrative presentations. Rapidly building a common base of understanding seems crucial in arriving at a point where fruitful discussions can occur regarding the impact of new technology, process streamlining, effects of regulations, and so forth.[14]

Once a model has been constructed, its usefulness can extend to other domains. For example, a model created to describe the factors that impact the success of the requirements engineering process[37] was subsequently applied to process design, customization, and reuse.[38]

Process improvement

The second goal cluster (Table 3) is concerned with evolutionary improvements to a process. A variety of reasons why an existing process might need to be changed include potential gains in pro-

ductivity, potential gains in product quality, external factors (such as schedule pressures or political factors) that require the process to be modified or steps to be removed, process refinement by the development team, and invalid assumptions under which the process was designed.

Software process models support the identification and analysis of potential improvements. This is a primary objective of software process modeling because of its connection with the improvement of software quality, cost, and scheduling. The models help highlight areas of opportunity for process improvements and can be used to evaluate potential improvements before they are put into practice. Models serve as storehouses for modifications, lessons learned, and tailoring decisions, thus recording the evolution of a process along with the outcomes of all changes made. This retrospection provides a basis for evaluating the relative success of such changes. In addition, tailoring decisions may be formalized and stored as part of the models, so that this knowledge can be consistently applied again in the future.[35]

There are many different approaches to software process improvement. A brief list includes the Quality Improvement Paradigm (QIP),[39] Total Quality Management (TQM),[40] and the SEPG application[41] from the SEI Capability Maturity Model.[4] Tool insertion in the software process is another approach that shows great promise.[42] In the next few sections, we describe the various research of PRS dedicated to process improvement.

Tool insertion method. Bruckhaus[43] describes the tool insertion method (TIM) as a comprehensive method that plans, executes, and controls the tool insertion into a software process. Currently, the method is mostly ad hoc. Tools are sometimes purchased based on informal recommendations and are put immediately into practice; often a tool does not live up to its expectations. Little or no planning is done for the use of a new tool in a particular process. It is often not clear what benefits can be expected when a specific tool is inserted into a specific process, nor how the impact of inserting the tool can be measured. Huff[44] describes some additional complexities of tool insertion:

- The actual cost of tool insertion may run five to eight times greater than the initial cost of the tool.
- It may take one to two years of preparation by a dedicated team of six to nine people to reach the pilot stage.
- Unanticipated costs may lead management to terminate a promising tool project or may increase resistance to future tool acquisitions.

Bruckhaus is currently developing TIM as part of a study to measure and analyze the impact of inserting a specific tool into an on-going, real-world software development process at the IBM Toronto Software Solutions Laboratory. As steps of TIM are executed, feedback is used to deliver improved versions of TIM. Software process modeling is the primary vehicle for quantifying expected benefits of, and measuring the impact of tool insertion.[45] TIM has eight steps:

1. Model the development activity of interest (if the model does not already exist)
2. Identify a set of key improvement areas
3. Measure current process context, data essence, and process performance
4. Quantify the goal of inserting a tool in terms of process performance
5. Select a tool and perform a pilot study
6. Customize the tool (if possible) to best suit the insertion goals
7. Insert the tool and monitor process performance
8. Take corrective actions (iterate)

TIM uses its own measurement method, TIM/M, to quantify the expected impact of tool insertion, and then monitors the actual tool usage to verify the forecasts. TIM/M is a refinement of Basili's goal-quality metric (GQM) paradigm,[46] with a specific focus on tool insertion. Like GQM, TIM/M defines a hierarchical structure of process aspects of interest. Measurement programs defined with the help of TIM/M cover the process aspects of process context, data essence, and process performance. Each of these aspects is then repeatedly broken down into subaspects. Finally, each aspect is coupled to one or more metrics that help describe the aspect. Since this structure is not limited to a particular process, TIM/M can help one design measurements in any given process context (for example, requirements planning, design, or testing). TIM will have its greatest impact on Goal 5 in Table 3.

The Organization and Process Together (OPT) approach. Seaman[47] describes Organization and

Process Together (OPT) as an approach to improve an organizational structure and a software development process when considered together as a single system. This method is patterned after QIP[39] and models the relationships between an organizational structure and a development process so that these relationships can be quantitatively measured and specific improvements can be planned in the system.

Relationships between the organizational structure and a development process are modeled using the Organizational and Architecture Specification Languages (OSL and ASL)[47] languages. An OSL specification describes the nonprocess relationships, called links, between different elements of the organization. Examples of such links are REPORTS-TO and MANAGES. An ASL specification describes the agents that execute the process in terms of the specific activities that they perform and the interactions between them. Since a process model often contains extreme detail, it is often difficult to model an organizational structure because the important information must be abstracted to be clearly identified. An ASL specification provides this abstraction and is more useful than a process model for analyzing the structure of an organization. It also provides the bridge between a development process and an organizational structure, allowing the two to be considered together as a system.

The goal of OPT is to measure how well suited a development process and an organizational structure are for each other. OPT targets two complementary factors that characterize the relationships between an organizational structure and a process. The insight OPT provides is that these factors can be measured. The first factor is the distribution of responsibility for process activities among the members of the organization. This captures the effect that the process has on the organization. The second factor is the process communication (which can either facilitate or hinder the efficient flow of information) within the organization. This captures the effect that the organization has on the process.

OPT is part of the planning phase (third step) of a QIP improvement cycle. The first QIP step results in the characterization of the environment by the creation of an organizational structure and a corresponding process model. After the initial organizational model is complete, the second QIP step is issued to create quantifiable goals using the GQM[46] approach. These goals define the means by which improvements will be validated. Reasonable expectations are calculated from the baseline provided by the first step. The third QIP step helps to decide the actions to take to satisfy the goals; the OPT approach is used for this purpose. As the actions are executed, the metrics chosen earlier are evaluated and analyzed to provide real-time feedback. Once all actions are completed, a postmortem analysis is conducted to determine if the original project goals have been satisfied. Any changes are integrated into the organizational model, and the cycle is prepared for the next QIP iteration. OPT addresses Goals 1 and 6, in Table 3, by validating a particular software process for a particular organizational structure. The key insight of OPT is that a process is not conducted in a vacuum; it affects, and is affected by, the organization that implements the process.

Software process models are used to predict the impact of changes.

Quantitative approaches. Raffo[48] describes a quantitative approach to process improvement. In a process change study, Raffo used software process models to predict the impact of potential process changes before a substantial commitment of time and resources was made. Such analysis allows process improvements to be prioritized based on their potential performance impact. This approach forecasts the impact, in quantitative terms, of a proposed process change before it is put into place in the actual organization. Using this method, management can ask what the value (impact) will be in the organization of making a (proposed) change to the process. This impact is measured in quantitative terms, such as effort (aggregate or time-profile) and schedule (total duration or intermediate milestones). This notion of impact focuses on the results of the change after it has been stabilized in the organization.

This investigation targets a real issue facing moderately mature organizations. As an example, one product team at the IBM Toronto Software Solutions Laboratory collected in a database about 200 proposals for improving processes employed in supporting a particular large product. The project manager then asked the process analyst on the project to prioritize these proposals on a cost and benefit basis. The process analyst felt reasonably able to estimate the implementation costs (such as documentation, training, and so forth), but had a need for a method of estimating the impact of a change in quantitative terms. The process change study directly addresses that need and helps develop a business case to obtain managerial support for specific process improvement proposals.

Raffo's approach uses simulation models of the affected portions of process with and without the proposed change. Stochastic modeling with Monte Carlo simulations is used, although deterministic modeling can also be performed (simpler, but less realistic). This approach explicitly models the complex interdependencies among process components and employs a new technique, called Task Element Decomposition, for handling interdependent operation times in large-scale systems; this is a contribution to the operations management field. A multi-attribute decision-making framework is used for comparing process alternatives, allowing management to ask questions about whether source code inspections would still be beneficial if the starting code had more errors, or would process performance be affected if slack time in the process were reduced.

The actual modeling techniques are built upon the SEI-developed Statemate-based modeling approach; this work most directly builds upon the modeling extensions to support quantitative management planning and control.[17] These software process modeling techniques also support sensitivity analyses of the models, helping the modeler:

- Assess confidence in the results
- Fine-tune the proposed process
- Explore learning curve effects
- Suggest viable alternative processes

The results to date are fully documented in a university working paper[49] totaling over one hundred pages. The first major phase of the work developed an approach to the quantitative comparison of alternative processes, together with supporting techniques—thus directly supporting Goal 3 in Table 3. A comprehensive representative case has been explored in detail, namely adding a code inspection process to an existing process that employed design reviews and unit tests, but not code inspections. The primary performance measures of interest were aggregate effort, duration, and defects remaining at the conclusion of the process. In this investigation, the process change resulted in somewhat more total effort, more total duration, but substantially fewer remaining errors.[49]

The long-term goal of this work is to develop a method and supporting techniques for forecasting the impact of potential process changes, thus directly supporting Goal 4 in Table 3. This work provides a vital foundation by developing an approach and technique for comparing alternative processes based on estimated quantitative performance from predictive software process models. A very preliminary method for the larger problem of process change has also been documented, but it requires further exploration. Additional example process changes will be considered, from which a method will be refined and documented and then tested on a full-scale real-world situation. This work is expected to be applicable to process changes at the scale of a Capability Maturity Model Key Process Area[5,6] or smaller, but may not be well suited to larger scale questions, such as comparing entire maturity levels.

Process generalization. Within a large organization, process improvement issues span multiple projects. There is a need, therefore, to improve not only individual processes in various projects (see sector B in Figure 1) but also generic process capabilities across a set of projects (sector A). The advantages of generic models include corporate-wide improvement and standardization of process components, development of an organization-specific culture (Goal 6 in Table 3), and a base model that may be tailored to meet the needs of specific projects.

Experience shows that while many organizations do indeed perform changes to prescriptive generic process models, they often have little or no information about the aspects of project-specific process models that need to be considered when changing the generic model. What is needed is a

mechanism that builds descriptive generic process models to provide a coherent vision of commonality and variability among project-specific processes. These descriptive generic models are useful for driving the desired changes to prescriptive generic models. At McGill, researchers have developed a method, called Generalizer, for building descriptive generic process models from a set of project-specific ones.[50] The key steps to this method are:

- Elicit project-specific process models
- Decompose and categorize project-specific models
- Match process components from different projects
- Obtain goal-oriented views from categorized components
- Make cross-project comparisons based on process views (Goal 3 in Table 3)
- Identify commonality among these views
- Synthesize common components into descriptive generic models

In its simplest form, a generalization method leads to a pure generic descriptive model. In practice, however, while making organization-wide change decisions, it is often desirable to examine various degrees of generic descriptive models as possible starting points for the changes. A suitable starting point is selected based on change-related, cultural, cost, technological, quality, standardization, or other factors. Using threshold values, process engineers can specify degrees of generality when building generic descriptive models. For example, a process engineer can set the threshold to 70 percent, implying that the descriptive generic model will contain all process components that are common to at least 7 out of 10 projects.

The output of this generalization method is descriptive generic models that show how project-specific models can be created from the organization's generic process models. A tool (named the Generalizer) is being developed to aid the building of generic prescriptive process models. It allows process engineers to experiment with threshold values to create the most suitable generic descriptive models. The Generalizer works in conjunction with the Elicit method and tool in that Elicit helps build descriptive process models and Generalizer helps generalize project-specific models.

Evolving processes. The approaches outlined so far can target specific improvements to a process model, but we need an additional mechanism that aids the evolution of an enacted process (Goal 7 in Table 3); the MARVEL project has such a tool, called the Evolver.[51] MARVEL[18] is a process-centered environment that defines a process model by a set of production system-style rules. Each process model is augmented by an object-oriented data schema that defines the process state and the composition of an object base that maintains all the data used by the process. As MARVEL enacts the process, it updates the process state stored in the data schema of the objects to accurately reflect its real-world state. Changes to a process model may directly affect existing objects. The process state, in particular, is maintained by attributes associated with each object that may need to be updated to be consistent with the new process model. Even some of the most trivial process changes might require the entire object base to be updated. For example, since each object inherits from a generic ENTITY class, adding an attribute to this class will update all objects.

There is a need to improve generic process capabilities.

The Evolver operates on an existing MARVEL environment (with a particular process model) and the evolution consists of two steps. First, the new data schema is compared against the original one, and a detailed analysis of their differences is reported. This allows the process engineer to view the consequences of a particular schema change. The second step involves the process model. MARVEL allows a process engineer to construct a sequence of rules whose forward chain must be atomically carried out; that is, if a concurrency conflict occurs at any point during the rule chain, the entire rule chain is rolled back. During this second step, the Evolver constructs a

graph from the original process model with rules as nodes and chains between rules as edges in the graph. A similar graph is constructed from the new process model and the two are compared to detect when these atomic rule sequences are shortened or extended. In the latter case, the Evolver automatically generates a batch script of MARVEL commands that executes the necessary rules in the new process model to make the object base consistent with respect to the new graph. In both cases, the process engineer is notified of all changes. The reports generated by the Evolver can be used to view the effects a particular process change will have (Goal 4 in Table 3); if the process engineer determines the change to be too costly, the evolution process can be terminated, restoring the MARVEL environment. This process evolution operates in an off-line fashion. Thus, the process itself can be in progress when evolution occurs but must be quiescent; that is, all atomic rule chains have terminated and the environment is waiting for the next request to continue the process. Since the main goal of the Evolver is its ability to resume long-lived processes after changes, this off-line approach is acceptable.

Achievements in improvement. A prerequisite to improving software processes is the need to understand them; such an understanding can be expressed in the form of models, as we have seen in the section "Human understanding and communication." Another important prerequisite is measurement. To facilitate measurement, valid and reliable instruments to measure various attributes of the process, and the software products, must be constructed. Recent work[52] has led to a method for the development of such instruments. This method uses software process models to define relevant metrics.[45]

Armed with a suitable set of metrics, and an understanding of processes, the improvement task has a higher likelihood of succeeding. Such an empirical perspective has been advocated as a basis for improvement.[53] A major benefit resulting from formalizing existing processes has been a substantial number of recommendations for process improvement. Many of the SEI-developed Statemate models have been subjected to analysis (manual and some automated) and recommendations were made for process improvement as well as for technology insertion. Some of the procedural issues observed from analyzing these models include bottlenecks, insufficient parallelism, and management control versus productivity. The opportunities for beneficial technology insertion include the use of advanced document production technology, configuration management tools, and project management and planning tools.[35]

Two examples, both from SEI, show the success of using process modeling techniques to guide process improvement. The first is the case of the process used by the U.S. Navy to support the operational software for the F-14A aircraft. The corresponding model depicted the full software support process, from receipt of a software trouble report, change request, or enhancement request, through release of the corresponding software change to the field. An extensive analysis of the model identified a total of 13 major issues for possible process improvements, resulting in over 30 recommendations for modifications to methods, procedures, and technology usage.[35] The second example is provided by the case of the military handbook (MIL-HDBK-347) on software support for mission-critical systems.[54] The SEI modeling and analysis identified 35 issues that could improve the handbook process or its exposition. These improvements were revealed after the completion of an extensive public review process that had solicited written comments from over two hundred people.[34]

The issue of change management, which is at the heart of Goal 7 in Table 3, has been investigated, where a model of changes, together with its infrastructure support, are presented in Reference 55. In addition to the provision of a framework for changes, a primary contribution of this work is the identification of the items of change and of their properties. Another look at the process evolution issue is provided in Reference 56, where the process cycle is used as a basis for process maintenance.

The experience base previously discussed is the centerpiece for any strategy to reuse software processes (Goal 2 in Table 3). Some results toward this direction can be found in the P/MARVEL environment.[57] This MARVEL environment is really a meta-process that allows a process engineer to design a process and verify its behavior in an isolated test environment. Multiple processes can be developed simultaneously, allowing process fragments to be transferred and

Table 4 Goals for automation (a combination of the original automated execution support and automated guidance)

	Goal
1	*Define an effective software development environment*
2	*Provide guidance, suggestions, and reference material to facilitate human performance of intended processes*
3	*Retain reusable process representations in a repository*
4	*Automate portions of processes*
5	*Support cooperative work among individuals and teams by automating process details*
6	*Automatically collect measurement data reflecting actual experience with a process*
7	*Enforce rules to ensure process integrity*

reused. P/MARVEL is further discussed in the next section.

Automation

The objectives of automated development support and automated guidance, shown in Table 4, have captured the interest of a large number of researchers. Early research on automation focused on providing a collection of independent file-based tools such as MAKE, VI, RCS, and SCCS. The invocation of these tools was, however, left up to the user. Software development environments (SDEs) then appeared that were considered "intelligent" since they automated some of the tool invocations and provided a means for storing information, such as the current state of a project, in databases. This intelligence was limited, however, because it was mostly fixed for each environment. SDEs were then created that could tailor their behavior based upon the specification of a desired process. These process-centered environments (PCEs) have process engines that *enact* process models, a term used to encompass enforcement, automation, and guidance of the users in carrying out the process.

Many existing process-centered environments use some form of rules to define software processes, because declarative rules are believed by many researchers to be the most natural way to express at least the local constraints on process steps. Major exceptions include Arcadia,[58] which uses an imperative notation called APPL/A[59] based on Ada; HFSP (hierarchical and functional software process),[60] which uses an extension of attribute grammars; and MELMAC,[61] SLANG,[62] and PROCESS WEAVER**,[63] which use a form of Petri nets. As previously described, MARVEL plays a central role in the PRS framework through its ability to enact processes. Before presenting MARVEL in more detail, we briefly describe some related PCEs.

Important examples of rule-based PCEs include GRAPPLE,[64] which applies a planning system to rules similar in form to those of MARVEL, and Darwin,[65] which employs backward chaining in the style of Prolog to enforce a set of "laws" that govern development activities and software changes. Most of the rule-based PCE projects are currently working toward support for multiple users, notably the Common Lisp Framework (CLF),[66] Oikos,[67] and Merlin.[68] CLF supports a checkout and merging model but has no central object base or process engine. Oikos uses a blackboard to communicate among separate workspaces, so there is somewhat more coordination required than with CLF. The approach Merlin uses is similar to that of MARVEL. To determine

whether a PCE is effective (Goal 1 in Table 4), we measure it against the set of requirements defined in Reference 18 which must be fulfilled by any general PCE. As described in that paper, MARVEL addresses each of the requirements.

MARVEL. The goal of the MARVEL[18,69] project is to develop a PCE kernel that guides and assists a team of users working on a medium-scale software development effort. The behavior of the generic kernel is tailored by an administrator who provides the schema, process model, tool envelopes, and coordination model for a specific project. The user, in contrast, generally sees only the resulting environment instance.

A process administrator writes a specification of the project data schema and process model using MSL, the process modeling language of MARVEL. The administrator then loads these specifications into the kernel, creating a MARVEL environment instance that supports both the data and process management requirements of the project. The data schema is specified in terms of classes, each of which consists of a set of typed attributes. Existing source code can be immigrated from the file system into a MARVEL object base using the "Marvelizer" utility.[70]

The administrator defines the process (or workflow) by creating process steps corresponding to individual software development tasks. Each step is encapsulated by a rule with a name and typed parameters. The body of a rule consists of a query to bind local variables, a complex logical condition on the actual parameters and bound variables that must be satisfied prior to initiating the activity of the step, an optional activity in which a software development tool may be invoked, and a set of effects, each of which asserts one of the activity's possible results (if there is no activity, there can be only one effect). Forward and backward chaining over the rules enforce consistency in the object base and automate tool invocations. Enforcement and automation are the two forms of enaction in MARVEL. This consistency enforcement is exactly the mechanism needed to satisfy Goal 7 in Table 4. The chains between rules form a rule network, with rules as nodes, and chains between rules as edges.

Process enaction is mainly user-driven, as opposed to system-driven. The user decides when to request a particular process step and enters a command with the name and actual parameters of the step. MARVEL then selects the "closest" matching rules (there may be more than one) and evaluates each of these rules in turn until it finds one whose condition is already satisfied or can be satisfied through backward chaining. The activity, if any, of this rule is then executed. Afterwards, one of the effects is selected according to a status code returned by the activity, and MARVEL forward chains to any other rules that are implications of this effect. If none of the conditions of the matching rules can be satisfied, however, then the user is informed that it is not possible to undertake that process step. Note that since rules have multiple effects, it may be possible that an attempted backward chain results in an undesired effect, but the chain is not then "undone" because that would be counter-productive (consider a backward chain to generate correct object code by compiling source code that produces syntax error messages instead). Additional details about the rule formalism and its chaining engine are given in Reference 69.

Multiple users of the same environment instance are supported by a client/server architecture.[18] A client provides the user interface, checks the arguments of commands, and executes tool envelopes; the process engine, synchronization management, and object base reside in the central MARVEL server. Scheduling is first-come, first-served, with rule chains interleaved at the natural breaks provided when clients execute activities. Clients may run on the same or different hosts as the server, but the enveloping facility assumes a shared network file system where the software components under development reside. The external view is illustrated in Figure 3. Additional details about multiuser issues, primarily concurrency control policies specified by the administrator in the coordination model, are found in References 71 and 72. The support for schema and process evolution, previously discussed, is described in Reference 51.

Synchronization among multiple users has three layers. Conventional locking is augmented by a lock compatibility matrix, part of the coordination model provided by the process administrator. This matrix provides support for composite objects by an ancestor lock table—a generalization of intention locks. Lock modes for kernel operations (for example, ADD, DELETE), as well as defaults for rule subparts and tool invocations,

Figure 3 Generic MARVEL environment

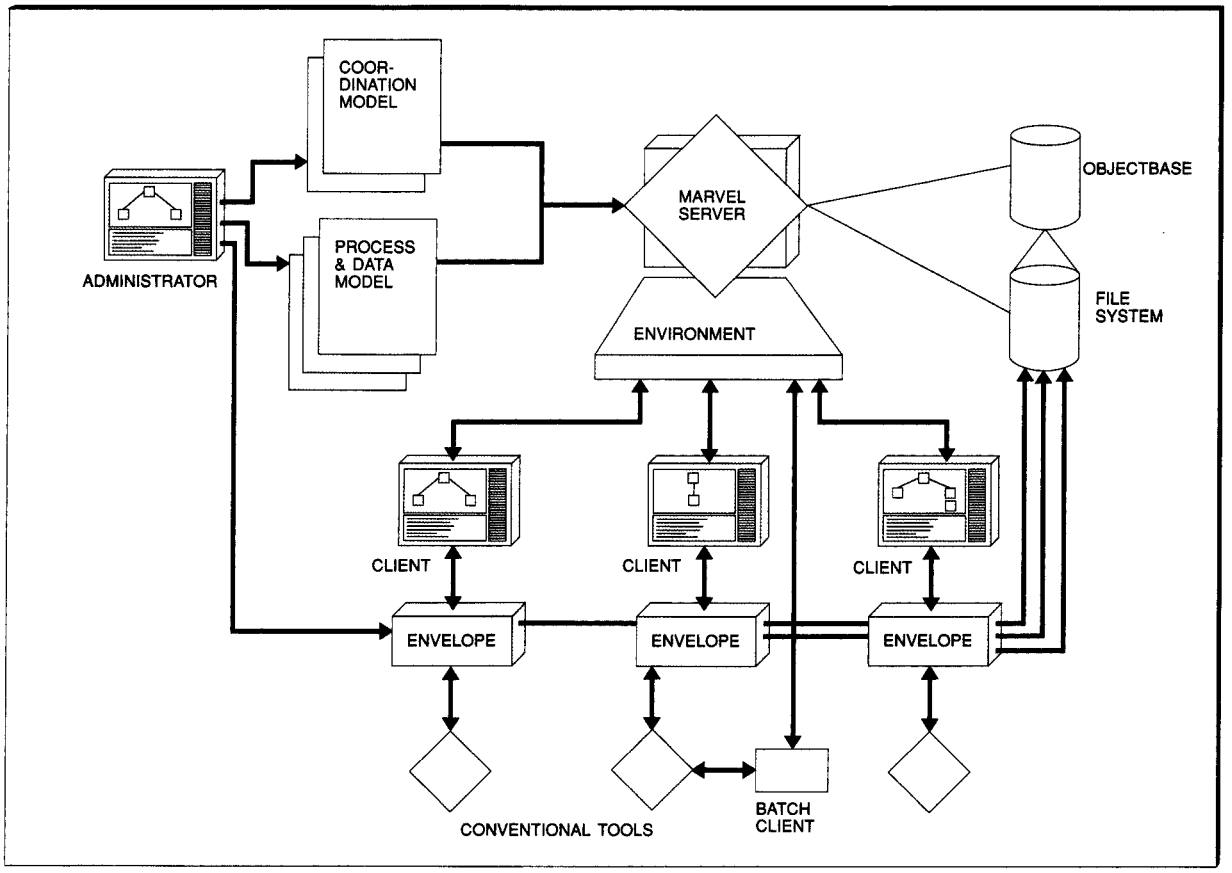

are also specified. The default concurrency control policy distinguishes between chaining for consistency versus automation purposes.[73] Chains for the purpose of maintaining consistency are mandatory and are treated as atomic, serializable transactions; if a consistency chain encounters an unresolvable lock conflict, the entire chain is aborted (rolled back). In contrast, chains for automation purposes are optional and are treated as sequences of distinct transactions, one for each rule; they can be terminated (and not rolled back) at rule boundaries. A preliminary coordination modeling language specifies scenarios where this default policy can be relaxed to increase concurrency and enhance collaboration.[71] The administrator defines the conflict resolution using primitive operations to NOTIFY a user, ABORT a rule chain, SUSPEND a rule chain until another has completed, or IGNORE the conflict.

Conventional file-oriented tools are integrated into a MARVEL process without source modifications, or even recompilation, through an enveloping language.[74] The rule activity indicates the tool and envelope name, with input literals and attributes to be supplied as arguments as well as output variables for binding to any returned results; an implicit status code selects the actual effect from among those given in the rule. The body of an envelope is a shell script, written in one of the conventional UNIX** shell languages: SH, KSH, or CSH.

Achievements in automation. Computer-assisted software engineering (CASE) tools were heralded as the solution to the software crisis when they were first developed; in retrospect, we have not seen this to be true. They have, however, proved to be extremely useful. The MARVEL PCE de-

Figure 4 Sample MARVEL rule network

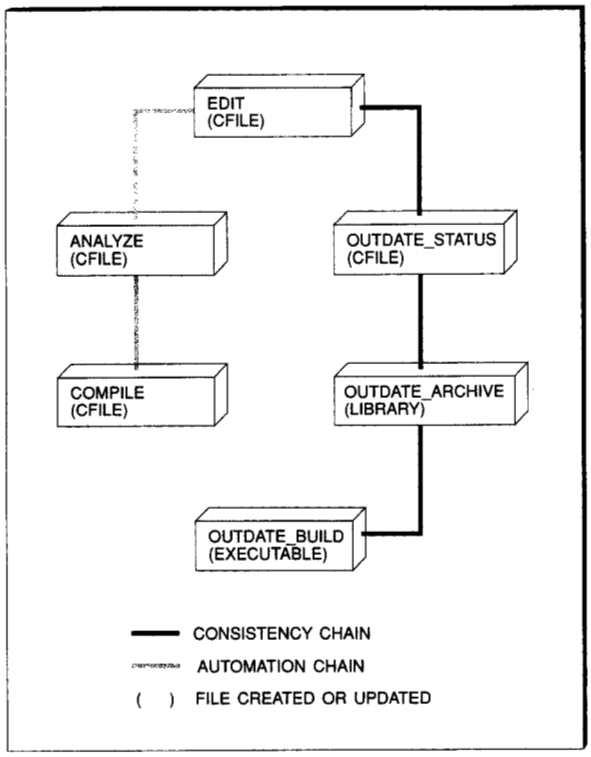

scribed here shows that processes can be, to some degree, automated to provide assistance to the users. MARVEL has been used to enact a variety of processes. The MARVEL code production environments (Oz/MARVEL[57] and C/MARVEL[18]) support the development of software using C. These environments manage code bases of 200 000 and 150 000 lines of code, respectively, and support teams of programmers. The Oz/MARVEL process was designed within our process development environment, P/MARVEL,[57] by taking the C/MARVEL process and reusing process fragments, tailoring, and adding new features. P/MARVEL allows multiple processes to be developed simultaneously and is a limited attempt at constructing a process repository (Goal 3 in Table 4). Doc/MARVEL is a document preparation environment using LaTeX, which the MARVEL group used to produce a four-volume set of manuals totaling over four hundred pages. In this environment, as many as five technical writers were working concurrently.

To provide guidance and reference material to the user (Goal 2 in Table 4), the MARVEL client allows the user to graphically browse the rule network to see the process flow. For example, Figure 4 contains a fragment of a rule network that shows what could happen if the user activates the edit rule on an object from the CFILE class. There is an atomic consistency chain of three rules that will be created if the user makes a change to the CFILE object. Once this chain has completed, an automation chain of two rules occurs, to analyze and compile the CFILE object.

The primary goal of MARVEL is to enact a process model (Goal 4 in Table 4); it does so in a user-driven fashion. That is, when the user completes a step of the process (by activating a rule), MARVEL uses forward chaining to carry out other process steps where prerequisites have become satisfied. MARVEL also executes backward chaining when a user initiates a process step where the prerequisite is not satisfied. In this case, the goal-directed backward chain attempts to satisfy the prerequisite by finding other process steps that, if executed, would make the original prerequisite satisfied.

MARVEL has no built-in mechanism to collect statistics on the actual experience of a process, but the administrator can design tool envelopes that record information about the process (Rule 6 in Table 4). For example, the Oz/MARVEL process has been written so that information is recorded every time the edit rule is activated on an object or a new version of an object is released by the configuration management system. Using this information, for example, the administrator can determine "hot spots" in the code where frequent changes are made. The enacted process can record as many statistics about itself as the administrator requires.

Cooperative work among small- to medium-sized groups of individuals is supported by MARVEL since it automates process-specific concurrency control policies while still supporting conventional transactions (that is, isolation while process fragments are in progress) as the default. MARVEL does not address collaboration among autonomous teams (Goal 5 in Table 4), but a new architecture for MARVEL is being created to support such collaboration among geographically dispersed teams across a wide area network such as the Internet.[75]

Table 5 Goals for process management

	Goal
1	*Develop a project-specific software process to accommodate the attributes of a particular project, such as its product, or organizational environment*
2	*Reason about attributes of software creation or evolution*
3	*Support development of plans for the project (forecasting)*
4	*Monitor, manage, and coordinate the process*
5	*Provide a basis for process measurement, such as definition of measurement points within the context of a specific process*

Process management

The final goal cluster (Table 5) addresses the planning, control, and operational management of software processes. Process management is a specific discipline that supports Total Quality Management (TQM)[40] and is needed to support continuous improvement of defined processes. Process management relies on feedback from measurements, assessments, and other analyses to guide improvement activities. Process management must also nurture an organizational culture that subscribes to a process-driven approach to software engineering and a continuous, on-going improvement.

The first step to process management is the selection of software processes most appropriate for the individual needs of a project and its organizational structure. In an earlier section we showed how OPT helps to determine the match between a particular process and the company's organization. Boehm[76,77] describes a "software process model generator" to aid in selecting the type of process for a given software effort based on a decision table and various life-cycle processes. Another important element of planning involves the development of a suitable process from a repertoire of components. The artificial intelligence planning paradigm used by GRAPPLE[64] provides an automated goal-driven approach to this problem. In GRAPPLE, process components are selected from an existing set and organized automatically to satisfy stated goals. The SFINX software process model library[78] offers a hierarchical organization of processes into layers, allowing components to be "plugged" into "frames" provided by the next higher layer of the architecture. Other works, such as those in References 34 and 79, explore reuse-based mechanisms for developing project-specific software processes.

Once a generic process model is selected it needs to be tailored to fit individual projects. Evolutionary approaches to process design have been characterized[34,80] and the two major categories are derivative and constructive. Further characterizations are based upon the source of the primary starting process and of the changes being integrated (as improvements, for tailoring or customization).

Process models can be useful in suggesting measurement points and metrics that can be used as status criteria (Goal 5 in Table 5). Indeed, it is valuable for a model to include definitions of precisely what is to be measured and when, to whom it is to be reported, how it is to be used, and so forth.[45,81] Such data should be collected and retained over time to reflect past experience and be used as a foundation for future planning.[35] One problem with measurements in a process model is that the actual performance of a process might differ from its model. Consequently, it is important to record the performance of the process (for

example, using a process trace), and consider all measurements against that trace. The history of past processes and results can be stored in a process database.[6]

Software process models can be used for planning schedules and analysis. References 15 and 17 describe modeling support for management planning and control (including monitoring, recording progress, and replanning), examining both general needs and specific capabilities available with Statemate-based modeling. Support for both planning and replanning during process performance was found to be an essential capability for managerial support. This support should enable planning schedules, costs, and resource needs with and without resource constraints, and accommodate deterministic or stochastic information.

Statemate-based modeling highlights importance of feedback loops.

The Statemate-based modeling approach is extended[17] to incorporate automated, quantitative simulations that are used to derive schedules, required work effort, and required staffing profiles. Cases of both point estimates (deterministic modeling) and uncertain estimates (stochastic modeling) are discussed, and resource constraints are also considered. This modeling approach offers the distinct advantage of smooth integration of representation, analysis, and forecasting capabilities; the quantitative simulations can be run by adding relatively straightforward information to an existing model, with no need to otherwise modify the existing model and with no changes to the visual representation. This integration is important, since this approach has been successfully used to model and analyze various large-scale real-world software processes.

The Statemate-based modeling approach also offers distinct advantages over traditional project management approaches such as the critical path method and PERT. The process models are more general, provide enhanced visibility into behavior, and highlight the importance of feedback loops in software processes. Moreover, they are amenable to resource constraints and full Monte Carlo simulation analysis.[17] Akhavi and Wilson have reported practical applications of these techniques at Rockwell International.[82]

Other work has also addressed some of these needs: SPMS, a software process modeling system,[83] contains a project management tool using the critical path method for project planning; system dynamics have been productively employed to forecast project level plans and the impact of changes.[84,85] Articulator[33] uses artificial intelligence scheduling techniques from production systems. Prism[11] incorporates process simulations.

Key process areas. A major component of SEI research at CAS is a study that examines the relationships and trade-offs among major dimensions of software total quality: life-cycle cost, field defects, and customer satisfaction. Various management-level questions are addressed, such as how the cost and satisfaction vary with defects, and how the defects and cost vary with front-end investment during development. Moreover, the research explores the factors underlying the differences and interrelationships seen. This work attempts to determine the key drivers of software total quality and find how they impact these major dimensions. For example, the value and impact of the following are determined:

- CASE design tools, automated testing tools, peer reviews, configuration management, and methodology training
- Each of the key process areas (KPAs) defined in the SEI Capability Maturity Model[5,6]

In terms very specific to software processes, practical questions regarding the value and impact of software process differences are addressed, such as the value typically obtained in practice from implementing a given KPA or, based on the actual experiences of others, the value an organization can expect to achieve from implementing a given KPA. Answers can be very useful in enabling management to obtain rough esti-

mates of the types and magnitudes of the values they might achieve from putting selected KPAs into practice. KPAs can be prioritized by these estimates for strategic attention, based on anticipated value. This would be quite appropriate for an organization that does not yet possess deep insight and understanding into its current processes, and is therefore not in a position to use the detailed quantitative process modeling techniques earlier discussed.

It would be useful to determine typical values achieved from individual KPAs, even at a broad average level. However, such a broad average would be a relatively crude estimate to apply to any single, specific case. For the purposes of a single organization or project, a model that takes other important factors (for example, type of application, personnel capability, or size of product) into account would provide a more refined estimate of likely value to that specific organization.

Krishnan[86,87] is conducting a field study based on data collected from the IBM Toronto Software Solutions Laboratory. It is important to recognize that this is an empirical investigation of actual results achieved in an actual organization. This directly addresses a vital real-world concern of reporting typical actual experience, not just theoretical possibility or best-case performance. Econometric statistical techniques (such as multiple regression and ordered probit analyses) are being used to analyze the data and test hypotheses.

The underlying conceptual model is relatively straightforward. Instead of considering software quality as being limited solely to conventional error or defect counts, software total quality is based on "delighting the customers."[88] Second, this research hypothesizes that software total quality is determined by the interactions of product, people, technology, process, and environmental factors. These factor groupings comprise a "system" of drivers that determine total quality.[89] Finally, the quality resulting from the driver interactions is measured along certain selected dimensions, which in this work include life-cycle cost, field quality (based on defects), and customer satisfaction.

In Phase I of this effort, the focus has been on understanding the interrelationships and tradeoffs among the measured result dimensions. Some specific research questions addressed are:

- Is quality "free" in the context of software? (Crosby[90] claims that it is.)
- Does higher deployment of resources in the early phases of systems development pay off in terms of quality?
- How does product type affect the relationship between cost and quality?

Early Phase I results are as follows. The empirically observed relationship between cost and quality indicates that as quality increases (lower shipped defect density), long-run unit cost (life-cycle cost per thousand lines of code) decreases, but at a decreasing rate. This confirms Crosby's assertion, at least over the range of quality observed in the data set. The relationships among resource deployment strategy (higher or lower investment early in the life cycle), quality, and product size have also been examined. The results indicate that (1) higher front-end investment pays off in terms of higher quality, and (2) the payoff in quality due to higher front-end investment is more pronounced for larger products. These results suggest the investment (by managerial action) in the early stages of the life cycle, especially for larger products, which appears to lead to higher quality and correspondingly lower costs.

Phase II will broaden the analysis to include key drivers of software quality and their impact on the measured dimensions of software quality. The measured result dimensions will be the same as in Phase I. The drivers will be specific factors reflecting the important characteristics of process, people, technology, product, and environment. The process factors will indicate degree of conformance with each of the identified KPAs. The other factors will measure specific characteristics of the other four driver groupings; the COCOMO cost drivers[20] provide examples of the factors to be examined. Phase II will identify key driver factors, the empirical relationships among them, and the major dimensions of software quality.

From a managerial perspective, the Phase II model will allow a manager to examine the result dimensions and a proposed change (for example, to emphasize increasing field quality), and then gain insight into what changes in the driver factors would yield the desired results. Those factors that

Figure 5 Hierarchy of objectives

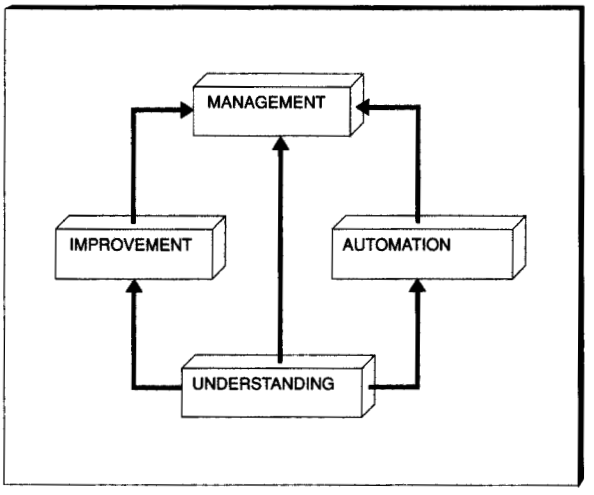

the manager can control (such as the process factors) can then be manipulated to obtain the desired results in practice. This illustrates how the work will yield managerial implications that can be acted upon.

The Phase II work will also yield broad estimates of value. The model itself will reveal the value of process (and people, technology, etc.) factors in terms of life-cycle cost, user quality, and customer satisfaction. Using average values for the nonprocess factors will yield the average value seen in the data set for each KPA (on an individual, incremental basis). For a specific project, one would estimate the values to be seen for the nonprocess factors, and the model could then be used to predict the value to be obtained for that project from each KPA.

Achievements in management. Software process models can enhance process management in a number of valuable ways, as discussed in this section. The feedback provided to management by measurements helps characterize current problems and areas of opportunity, and may also lead to further process improvements (Goal 3 in Table 5); thus this process management goal cluster also relates to the process improvement cluster. Better process management will also be greatly aided by accomplishing some of the other goals, such as increased understanding, better training, conformity to process definitions, and evaluation of potential improvements. The process cycle[10] and OPT[47] can be used together to create project-specific software processes (Goal 1 in Table 5) that are best suited for a particular product. MARVEL[18] addresses Goal 4 in Table 5 by allowing the process manager to observe the behavior of the process and define coordination policies for the multiple users.[71] MARVEL provides mechanisms that allow the administrator to improve and evolve the process as feedback from the performers is acquired.

Conclusion

The consortium study to improve the quality of software processes was initiated at a time when process awareness was very low. The awareness has matured as the research on software processes has increased. By providing a holistic vision of software processes, the results of the study can affect many aspects of processes, as we have seen in this paper. Table 6 summarizes the research results corresponding to the 24 subgoals. This table concisely shows the breadth of process research performed by the study participants.

After some reflection on the nature of these objectives, we can create a hierarchy as shown in Figure 5. The base objective, understanding, is a precursor to all other objectives. Improvement and automation are independent of each other, and reside at the same level in the hierarchy, while management, at the highest level, is dependent upon all other objectives. This hierarchy portrays the dependencies between the objectives and reflects the depth of our research. The presentation of the understanding objective in the section "Human understanding and communication," for example, is much more comprehensive than the discussion on management in the section "Process management." We plan to continue research addressing each of these four objectives, ultimately providing real-world solutions to real-world problems.

Finally, there are tangible benefits to IBM in the continuation of these studies.

1. There is interest within IBM in the application of the process technology. We are currently working with several development groups at IBM that are interested in applying process technology.
2. The IBM staff will learn about new process

Table 6 Summary of research results

Objective	Goal number	Results (Cited reference number in paper)
Understanding (Table 2)	1-5	Elicit (22, 23, 25), DMP SEI Statemate approach (7, 13, 14, 30, 34, 35)
Improvement (Table 3)	1	OPT (47)
	2	P/MARVEL, (34)
	3	(15, 17, 48, 49)
	4	(48, 49), Evolver (51)
	5	TIM (43), (13, 14, 35)
	6	OPT, (34, 86)
	7	Evolver, Prism (11), (7, 80)
Automation (Table 4)	1-7	MARVEL (18), OZ (75), P/MARVEL, (34)
Management (Table 5)	1	(34, 80), OPT
	2	Process Cycle (10), (7, 15, 35, 86, 87)
	3	(15, 17, 86, 87)
	4	MARVEL, (17)
	5	(45)

methods and techniques, new tools, and novel process models.

3. The IBM staff will increase its contact with the academic world, and thus be exposed to real-world problems and assist in the solutions to those problems.

Acknowledgments

The authors would like to thank the current PRS students for their continuing hard work and research: Tilmann Bruckhaus, M. S. Krishnan, David Raffo, and Carolyn B. Seaman. Previous CAS students involved with PRS include Dirk Höltje, Steven Popovich, and Andrew Z. Tong. Sincere thanks are also due to research students in projects other than PRS who have contributed greatly to our process work in general: Israel Z. Ben-Shaul, Khaled El Emam, Won-Kook Hong, Graciela Perez, Khalid Sherdil, Peter D. Skopp, Kamel Toubache, and Josee Turgeon. We would like to thank the CAS staff—in particular Jacob Slonim—for providing an open and challenging research environment.

**Trademark or registered trademark of i-Logix, Inc., Cap Gemini Innovation, or X/Open Co. Ltd.

Cited references and notes

1. K. D. Saracelli and K. F. Bandat, "Process Automation in Software Application Development," *IBM Systems Journal* **32,** No. 3, 376–396 (1993).
2. ISO 9000-3, *Guidelines for the Application of ISO 9001 to the Development, Supply and Maintenance of Software*, International Organization for Standardization, Geneva (1991).
3. D. A. Garvin, "How the Baldrige Award Really Works," *Harvard Business Review* **69,** No. 6, 80–93 (November–December 1991).
4. M. Paulk, B. Curtis, M. Chrissis, and C. Weber, "Capability Maturity Model, Version 1.1," *IEEE Software* (July 1993), pp. 18–27.
5. M. C. Paulk, B. Curtis, M. B. Chrissis, and C. V. Weber, *Capability Maturity Model for Software, Version 1.1*, Technical Report CMU/SEI-93-TR-24; Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA 15213 (February 1993).
6. M. C. Paulk, C. V. Weber, S. M. Garcia, M. B. Chrissis, and M. Bush, *Key Practices of the Capability Maturity Model, Version 1.1*, Technical Report CMU/SEI-93-TR-25; Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA 15213 (February 1993).
7. B. Curtis, M. I. Kellner, and J. W. Over, "Process Modeling," *Communications of the ACM* **35,** No. 9, 75–90 (September 1992).
8. P. H. Feiler and W. S. Humphrey, *Software Process Development and Enactment: Concepts and Definitions*, Technical Report SEI-92-TR-4, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA 15213 (September 1992).
9. W. S. Humphrey, *Managing the Software Process*, Addison-Wesley Publishing Co., Reading, MA (1989).
10. N. H. Madhavji, "The Process Cycle," *Software Engineering Journal* **6,** No. 5, 234–242 (September 1991).
11. N. H. Madhavji, "The Prism Model of Changes," *Proceedings of the 13th International Conference on Software Engineering*, Austin, TX, May 1991; IEEE Computer Society Press, pp. 93–96.
12. V. R. Basili, G. Caldiera, and G. Cantone, "A Reference Architecture for the Component Factory," *ACM Transactions on Software Engineering and Methodology* **1,** No. 1, 53–80 (January 1992).
13. M. I. Kellner and G. A. Hansen, *Software Process Modeling*, Technical Report CMU/SEI-88-TR-9, DTIC: ADA197137, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA 15213 (May 1988).
14. M. I. Kellner and G. A. Hansen, "Software Process Modeling: A Case Study," *Proceedings of the 22nd Annual Hawaii International Conference on System Sciences, Volume II*, Kona, HI, January 1989; IEEE Computer Society Press, pp. 175–188.
15. W. S. Humphrey and M. I. Kellner, "Software Process Modeling: Principles of Entity Process Models," *Proceedings of the 11th International Conference on Software Engineering*, Pittsburgh, PA, May 1989; IEEE Computer Society Press, pp. 331–342.
16. D. Harel et al., "Statemate: A Working Environment for the Development of Complex Reactive Systems," *IEEE Transactions on Software Engineering* **16,** No. 4, 403–414 (April 1990).
17. M. I. Kellner, "Software Process Modeling Support for Management Planning and Control," *Proceedings of the 1st International Conference on the Software Process: Manufacturing Complex Systems*, Redondo Beach, CA, October 1991, M. Dowson, Editor; IEEE Computer Society Press, pp. 8–28.
18. I. Z. Ben-Shaul, G. E. Kaiser, and G. T. Heineman, "An Architecture for Multi-User Software Development Environments," *Computing Systems* **6,** No. 2, 65–103 (Spring 1993).
19. G. T. Heineman, *Automatic Translation of Process Modeling Formalisms*, Technical Report CUCS-036-93, Department of Computer Science, Columbia University, New York (November 1993).
20. B. W. Boehm, *Software Engineering Economics*, Prentice-Hall, Inc., Englewood Cliffs, NJ (1981).
21. L. Osterweil, "Software Processes Are Software Too," *Proceedings of the 9th International Conference on Software Engineering*, IEEE Computer Society Press (March 1987), pp. 2–12.
22. N. H. Madhavji, W. K. Hong, T. Bruckhaus, and J. E. Botsford, *Elicit: A Meta Process and Supporting Tool for Eliciting Software Process Models*, Technical Report SE-92.4, McGill University, Montreal (September 1992).
23. N. H. Madhavji, D. Höltje, W. Hong, and T. Bruckhaus, *Eliciting a Formal Model of a Requirements Engineering Process*, Technical Report SE-93.1, School of Computer Science, McGill University, Montreal (August 1993).
24. N. H. Madhavji, D. Höltje, W. Hong, and T. Bruckhaus, *An Empirically Improved Process for Modelling Software Processes*, Technical Report SE-93.3, School of Computer Science, McGill University, Montreal (August 1993).
25. D. Höltje, *Building Software Process Models Using the Elicit Meta Process: A Case Study*, master's thesis, Fern Universität at Hagen, Germany. This research was carried out at McGill University, Montreal, and IBM Canada Ltd. Laboratory, Toronto (June 1993).
26. R. K. Keller and N. H. Madhavji, "A Comprehensive Process Model for Studying Software Process Papers," *Proceedings of the 15th International Conference on Software Engineering*, Baltimore, MD, May 1993; IEEE Computer Society Press, pp. 78–88.
27. N. H. Madhavji, K. Toubache, and W. Hong, "Towards Engineering Reliable Software Processes," *International Software Quality Exchange (ISQE92)*, San Francisco, CA, March 1992; Juran Institute, Inc., and the Rocky Mountain Institute of Software Engineering, pp. 4B-1-30.
28. N. H. Madhavji, K. Toubache, and W. Hong, "Communications and Iterations in the Process Cycle," *Proceedings of the 7th International Software Process Workshop*, Yountville, CA, October 1991; IEEE Computer Society Press, pp. 91–93.
29. N. H. Madhavji, K. Toubache, and E. Lynch, "The IBM-McGill Project on Software Process," *Proceedings of the 1991 CASCON Conference*, Toronto, October 1991; IBM Canada Ltd. Laboratory, Toronto, pp. 95–109.
30. M. I. Kellner, "Representation Formalisms for Software Process Modeling," *Proceedings of the 4th International Software Process Workshop: Representing and Enacting the Software Process*, C. Tully, Editor; *ACM Software Engineering Notes* (June 1989), pp. 93–96. Also special issue of *ACM Software Engineering Notes* **14,** No. 4.

31. N. H. Madhavji and W. Schäfer, "Prism—Methodology and Process-Oriented Environment," *IEEE Transactions on Software Engineering* **17,** No. 12, 1270–1283 (December 1991).
32. P. Mi and W. Scacchi, "A Knowledge-Based Environment for Modeling and Simulating Software Engineering Processes," *IEEE Transactions on Knowledge and Data Engineering* **2,** No. 3, 283–294 (September 1990).
33. P. Mi and W. Scacchi, "Modeling Articulation Work in Software Engineering Processes," *Proceedings of the 1st International Conference on the Software Process*, Redondo Beach, CA, October, 1991; IEEE Computer Society Press, pp. 188–201.
34. M. I. Kellner and R. W. Phillips, "Practical Technology for Process Assets," *Proceedings of the 8th International Software Process Workshop*, Wadern, Germany, March 1993, W. Schäfer, Editor, pp. 107–112.
35. M. I. Kellner, "Software Process Modeling: Value and Experience," *SEI Technical Review*, pp. 23–54; Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA 15213 (1989).
36. A United States Air Force term referring to technical documentation for end users.
37. K. El Emam and N. H. Madhavji, *A Study of the Factors That Affect the Success of the Requirements Engineering Process*, Technical Report SE-94.2, School of Computer Science, McGill University, Montreal (1994).
38. G. Perez, K. El Emam, and N. H. Madhavji, *A Method for the Evaluation of Congruence and Its Application to Software Process Design and Customization*, Technical Report SE-93.4; School of Computer Science, McGill University, Montreal (1993).
39. V. R. Basili and H. D. Rombach, "Tailoring the Software Process to Project Goals and Environments," *Proceedings of the 9th International Conference on Software Engineering*, Monterey, CA, April 1987; IEEE Computer Society Press, pp. 345–357.
40. J. Hauser and D. Clausing, "The House of Quality," *Harvard Business Review* **66,** No. 3, 63–73 (May–June 1988).
41. P. Fowler and S. Rifkin, *Software Engineering Process Group Guide*, Technical Report CMU/SEI-90-TR-24, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA 15213 (September 1990).
42. *Conference on Transferring Software Engineering Tool Technology*, Cat # 88TH0218-8 (November 1987); available from IEEE Service Center, Piscataway, NJ.
43. T. Bruckhaus, "The Impact of Inserting a Tool into a Software Process," *Proceedings of the 1993 CASCON Conference*, Toronto, October 1993, A. Gawman, W. M. Gentleman, E. Kidd, P.-A. Larson, and J. Slonim, Editors; IBM Canada Ltd. Laboratory, Toronto, and National Research Council, Canada, pp. 250–264.
44. C. C. Huff, "Elements of a Realistic Case Tool Adoption Budget," *Communications of the ACM* **35,** No. 4, 45–54 (April 1992).
45. N. H. Madhavji, J. Botsford, T. Bruckhaus, and K. El Emam, "Measurements Based on Process and Context Models," *Proceedings of the International Workshop on Experimental Software Engineering Issues*, Dagstuhl, Warden, Germany, September 1992, V. R. Basili, D. H. Rombach, and R. Selby, Editors; Springer-Verlag, pp. 67–72.
46. V. R. Basili and H. D. Rombach, "The TAME Project: Towards Improvement-Oriented Software Environments," *IEEE Transactions on Software Engineering* **14,** No. 6, 758–773 (June 1988).
47. C. B. Seaman, "OPT: Organization and Process Together," *Proceedings of the 1993 CASCON Conference*, Toronto, October 1993, A. Gawman, W. M. Gentleman, E. Kidd, P.-A. Larson, and J. Slonim, Editors; IBM Canada Ltd. Laboratory, Toronto, and National Research Council, Canada, pp. 314–324
48. D. M. Raffo, "Evaluating the Impact of Process Improvements Quantitatively Using Process Modeling," *Proceedings of the 1993 CASCON Conference*, Toronto, October 1993, A. Gawman, W. M. Gentleman, E. Kidd, P.-A. Larson, and J. Slonim, Editors; IBM Canada Ltd. Laboratory, Toronto, and National Research Council, Canada, pp. 290–313.
49. D. M. Raffo, *Assessing the Impact of Potential Process Changes for Large Scale Software Development Using Process Modeling*, Working Paper WP#1993-14, Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, PA 15213 (April 30, 1993).
50. W. Hong and N. H. Madhavji, *A Method for Building Generic Process Models*, Technical Report SE-94.1, School of Computer Science, McGill University, Montreal (January 1994).
51. G. E. Kaiser, I. Z. Ben-Shaul, G. T. Heineman, and W. Marrero, *Process Evolution for the Marvel Environment*, Technical Report CUCS-047-92, Department of Computer Science, Columbia University, New York (April 1993).
52. K. El Emam, N. Moukheiber, and N. H. Madhavji, "An Empirical Evaluation of the G/Q/M Method," *Proceedings of the 1993 CASCON Conference*, Toronto, October 1993, A. Gawman, W. M. Gentleman, E. Kidd, P.-A. Larson, and J. Slonim, Editors; IBM Canada Ltd. Laboratory, Toronto, and National Research Council, Canada, pp. 265–289.
53. K. El Emam, N. H. Madhavji, and K. Toubache, "Empirically Driven Improvements of Generic Process Models," *Proceedings of the 8th International Software Process Workshop*, Wadern, Germany, March 1993, W. Schäfer, Editor, pp. 61–65.
54. *Military Handbook: Mission-Critical Computer Resources Software Support*, MIL-HDBK-347, U.S. Department of Defense (May 1990).
55. N. H. Madhavji, "Environment Evolution: The Prism Model of Changes," *IEEE Transactions on Software Engineering* **18,** No. 5, 380–392 (May 1992).
56. N. H. Madhavji, K. Toubache, and W. Hong, "A Framework for Process Maintenance," *Proceedings of the IEEE 1992 Conference on Software Maintenance*, Orlando, FL, November 1992; IEEE Computer Society Press, pp. 245–254.
57. "Programming Systems Laboratory," *Marvel 3.1 Administrator's Manual*, Technical Report CUCS-009-93, Department of Computer Science, Columbia University, New York (March 1993).
58. R. Kadia, "Issues Encountered in Building a Flexible Software Development Environment," *Proceedings of the 5th ACM SIGSOFT Symposium on Software Development Environments*, Tyson's Corner, VA, December 1992, H. Weber, Editor, pp. 169–180. Also special issue of *ACM Software Engineering Notes* **17,** No. 5.
59. S. M. Sutton, Jr., *APPL/A: A Prototype Language for Software-Process Programming*, Ph.D. thesis, University of Colorado (August 1990).

60. T. Katayama, "A Hierarchical and Functional Software Process Description and Its Enaction," *Proceedings of the 11th International Conference on Software Engineering*, Pittsburgh, PA, May 1989; IEEE Computer Society Press, pp. 343–352.
61. W. Deiters and V. Gruhn, "Managing Software Processes in the Environment Melmac," *Proceedings of the 4th ACM SIGSOFT Symposium on Software Development Environments*, Irvine, CA, December 1990, R. N. Taylor, Editor, pp. 193–205. Also special issue of *ACM Software Engineering Notes* **15,** No. 6.
62. S. Bandinelli and A. Fuggetta, "Computational Reflection in Software Process Modeling: The SLANG Approach," *Proceedings of the 15th International Conference on Software Engineering*; IEEE Computer Society Press (May 1993), pp. 144–154.
63. C. Fernström, "PROCESS WEAVER: Adding Process Support to UNIX," *Proceedings of the 2nd International Conference on the Software Process: Continuous Software Process Improvement*, Berlin, February 1993; IEEE Computer Society Press, pp. 12–26.
64. K. E. Huff and V. R. Lesser, "A Plan-Based Intelligent Assistant That Supports the Software Development Process," *Proceedings of the ACM SIGSOFT/SIGPLAN Software Engineering Symposium on Practical Software Development Environments*, Boston, MA, November 1988, P. Henderson, Editor; ACM Press, New York, pp. 97–106. Also special issue of *ACM Software Engineering Notes* **13,** No. 5.
65. N. H. Minsky, "Law-Governed Systems," *Software Engineering Journal* **6,** No. 5, 285–302 (September 1991).
66. R. Balzer, "Tolerating Inconsistency," *Proceedings of the 13th International Conference on Software Engineering*, Austin, TX, May 1991; IEEE Computer Society Press, pp. 158–165.
67. V. Ambriola, P. Ciancarini, and C. Montangero, "Software Process Enactment in Oikos," *Proceedings of the 4th ACM SIGSOFT Symposium on Software Development Environments*, Irvine, CA, December 1990, R. N. Taylor, Editor, pp. 183–192. Also special issue of *ACM Software Engineering Notes* **15,** No. 6.
68. W. Schäfer, B. Peuschel, and S. Wolf, "A Knowledge-Based Software Development Environment Supporting Cooperative Work," *International Journal on Software Engineering & Knowledge Engineering* **2,** No. 1, 79–106 (March 1992).
69. G. T. Heineman, G. E. Kaiser, N. S. Barghouti, and I. Z. Ben-Shaul, "Rule Chaining in Marvel: Dynamic Binding of Parameters," *IEEE Expert* **7,** No. 6, 26–32 (December 1992).
70. M. H. Sokolsky and G. E. Kaiser, "A Framework for Immigrating Existing Software into New Software Development Environments," *Software Engineering Journal* **6,** No. 6, 435–453 (November 1991).
71. N. S. Barghouti, *Concurrency Control in Rule-Based Software Development Environments*, Ph.D. thesis, CUCS-001-92, Columbia University, New York (February 1992).
72. G. T. Heineman, *A Transaction Manager Component for Cooperative Transaction Models*, Ph.D. thesis proposal, CUCS-017-93, Department of Computer Science, Columbia University, New York (July 1993).
73. N. S. Barghouti, "Supporting Cooperation in the Marvel Process-Centered SDE," *Proceedings of the 5th ACM SIGSOFT Symposium on Software Development Environments*, Tyson's Corner, VA, December 1992, H. Weber, Editor, pp. 21–31. Also special issue of *ACM Software Engineering Notes* **17,** No. 5 (December 1992).
74. M. A. Gisi and G. E. Kaiser, "Extending a Tool Integration Language," *Proceedings of the 1st International Conference on the Software Process: Manufacturing Complex Systems*, Redondo Beach, CA, October 1991, M. Dowson, Editor; IEEE Computer Society Press, pp. 218–227.
75. I. Z. Ben-Shaul and G. E. Kaiser, "A Paradigm for Decentralized Process Modeling and Its Realization in the OZ Environment," *Proceedings of the 16th International Conference on Software Engineering*, Sorrento, Italy, May 1994; IEEE Computer Society Press, pp. 179–188.
76. B. Boehm and F. Belz, "Experience with the Spiral Model As a Process Model Generator," *Proceedings of the 5th International Software Process Workshop: Experience with Software Process Models*, Kennebunkport, ME, October 1989, D. Perry, Editor; IEEE Computer Society Press, pp. 43–45.
77. B. Boehm, "What We Really Need Are Process Model Generators," *Proceedings of the 11th International Conference on Software Engineering*, Pittsburgh, PA, May 1989; IEEE Computer Society Press, p. 397.
78. G. Bux and G. Marzano, "Library of Predefined Software Process Models As Support for Software Factory Design: The SFINX Proposal," *Proceedings of the 5th International Workshop on Computer-Aided Software Engineering*, Montreal, July 1992; IEEE Computer Society Press, pp. 176–178.
79. V. R. Basili and H. D. Rombach, "Support for Comprehensive Reuse," *Software Engineering Journal* **6,** No. 5, 303–316 (September 1991).
80. M. I. Kellner and L. P. Gates, "Evolution of Software Processes," *Proceedings of the International Workshop on the Evolution of Software Processes*, Mt. St. Hilaire, Quebec, January 1993.
81. H. D. Rombach and L. Mark, "Software Process & Product Specifications: A Basis for Generating Customized Software Engineering Information Bases," *Proceedings of the 22nd Annual Hawaii International Conference on System Sciences, Volume II*, Kona, HI, January 1989; IEEE Computer Society Press, pp. 165–174.
82. M. Akhavi and W. Wilson, "Dynamic Simulation of Software Process," *Proceedings of the 5th Software Engineering Process Group National Meeting*, Costa Mesa, CA, April 1993; Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA 15213 (1993). Presentation in Session Three, Advanced Track.
83. H. Krasner et al., "Lessons Learned from a Software Process Modeling System," *Communications of the ACM* **35,** No. 9, 91–100 (September 1992).
84. T. K. Abdel-Hamid and S. E. Madnick, *Software Project Dynamics: An Integrated Approach*, Prentice-Hall, Inc., Englewood Cliffs, NJ (1991).
85. C. Y. Lin and R. R. Levary, "Computer-Aided Software Development Process Design," *IEEE Transactions on Software Engineering* **15,** No. 9, 1025–1037 (September 1989).
86. M. S. Krishnan et al., "Cost, Quality, and User Satisfaction of Software Products: A Field Study," *Proceedings of the Field Studies in Quality Management Conference*, Simon School, University of Rochester, Rochester, NY, March 26–27, 1993.
87. M. S. Krishnan, "Cost, Quality, and User Satisfaction of

Software Products: An Empirical Analysis," *Proceedings of the 1993 CASCON Conference*, Toronto, October 1993, A. Gawman, W. M. Gentleman, E. Kidd, P.-A. Larson, and J. Slonim, Editors; IBM Canada Ltd. Laboratory, Toronto, and National Research Council, Canada, pp. 400–411.
88. The authors believe that this should be attributed to John Schwarz, IBM Toronto Software Solutions Laboratory.
89. M. I. Kellner and J. W. Over, "A Software Quality Improvement Framework," *Proceedings of the Software Engineering Symposium—1992*, Milan, June 10–11, 1992; Advanced Software Technology, Olivetti Information Services.
90. P. B. Crosby, *Quality Is Free*, McGraw-Hill, Inc., New York (1979).

Accepted for publication April 28, 1994.

George T. Heineman *Department of Computer Science, 450 CS Building, Columbia University, New York, New York 10027 (electronic mail: heineman@cs.columbia.edu).* Mr. Heineman is a Ph.D. candidate in the Computer Science Department at Columbia University. His research interests include cooperative transactions, software technology, and the intersection of process-centered environments with database technology. He received his B.A. degree in computer science from Dartmouth College and his M.S. degree from Columbia University. He is a member of the IEEE and the ACM.

John E. Botsford *IBM Canada Ltd., 844 Don Mills Road, North York, Ontario, Canada M3C 1V7 (electronic mail: bottsford@vnet.ibm.com).* Mr. Botsford is a research staff member at IBM Canada Ltd.'s Centre for Advanced Studies in the Toronto laboratory. He joined IBM in 1964 and worked in two branch offices before moving to the Toronto laboratory in late 1967. Since then he has worked on a large number of development projects and was an instructor in the education department for three years.

Gianluigi Caldiera *Institute for Advanced Computer Studies, Department of Computer Science, University of Maryland, College Park, Maryland 20742 (electronic mail: gcaldiera@cs.umd.edu).* Mr. Caldiera is on the faculty of the University of Maryland Institute for Advanced Computer Studies, where he coordinates projects on software quality and reuse. He is also associated with the Software Engineering Laboratory of NASA Goddard Space Flight Center, Greenbelt, Maryland. He has participated, as both project leader and project reviewer, with the European Strategic Program for Information Technology (ESPRIT). His research and professional activities are in software engineering, with special focus on software quality assurance and management, software metrics, software reuse, and software factories. He has worked on these topics with industry and government in both the U.S.A. and Europe. He has authored and coauthored many papers in international journals and conferences.

Gail E. Kaiser *Department of Computer Science, 450 CS Building, Columbia University, New York, New York 10027 (electronic mail: kaiser@cs.columbia.edu).* Dr. Kaiser is an Associate Professor of Computer Science and Director of the Programming Systems Laboratory at Columbia University. She was selected for an IBM Research Initiation Grant in Complex Information Systems in 1988. Dr. Kaiser has published over 80 papers in a range of areas, including software development environments, software processes, extended transaction models, object-oriented languages and databases, and parallel and distributed systems. Dr. Kaiser is an associate editor of the journal *ACM Transactions on Software Engineering and Methodology* and serves on numerous program committees for conferences. In addition she reviews papers for conferences, journals, NSF, and NSERC. She received her Ph.D. and M.S. from Carnegie Mellon University and her Sc.B. from the Massachusetts Institute of Technology. She is a member of AAAI and the ACM and a senior member of the IEEE.

Marc I. Kellner *Software Engineering Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213-3890 (electronic mail: mik@sei.cmu.edu).* Dr. Kellner is employed as a senior scientist at the Software Engineering Institute (SEI) at Carnegie Mellon University in Pittsburgh, Pennsylvania. He has pioneered and led much of the research on software process modeling conducted at the SEI, and has published several papers on software process issues. Prior to joining the SEI, Dr. Kellner was a professor at Carnegie Mellon University, where he established and directed a degree program in information systems. He received his Ph.D. in systems sciences (specializing in MIS) from the Graduate School of Industrial Administration at Carnegie Mellon University. His research interests include software processes, software process modeling, software maintenance, and quality management for software. He is a member of the ACM, the IEEE Computer Society, and The Institute of Management Sciences.

Nazim H. Madhavji *School of Computer Science, McGill University, McConnell Engineering Building, 3480 University Street, Montreal, Quebec, Canada H3A 2A7 (electronic mail: madhavji@opus.cs.mcgill.ca).* Dr. Madhavji obtained his Ph.D. degree in 1980 from the University of Manchester (U.K.). He joined McGill University, Montreal, Quebec, Canada in 1983, where he is a professor at the School of Computer Science. In 1993, he was appointed as Research Director of the Software Process Programme at *Centre de recherche informatique de Montréal* (CRIM). His research interests are in software engineering, software processes, project management, software environments, and programming languages. He leads ProM Canada, the Canadian component of a Canada-Germany joint research project in software processes involving McGill/CRIM, GMD (German National Research Centre in Computer Science) and Fern Universität, Hagen. He is a principal investigator in a reverse engineering project, involving McGill University, the University of Toronto, the University of Victoria, and IBM Canada. He has chaired, and cochaired, numerous program committees, workshops, and conference sessions. He is on the Advisory Editorial Board of the *Journal of Software Maintenance*. He has led Quebec and Canadian missions to foreign countries on the subject of software engineering, software processes, and CASE technologies. He is a consultant to several organizations in the field of software engineering and process technology.

Reprint Order No. G321-5553.