THE ARCADIA ENVIRONMENT ARCHITECTURE

Richard N. Taylor, Frank C. Belz† Lori A. Clarke*,
Leon Osterweil**, Richard W. Selby, Jack C. Wileden*,
Alexander L. Wolf*, Michal Young

COINS Technical Report 87-97
September 1987

Department of Information and Computer Science
University of California, Irvine[1]

*Department of Computer and Information Science
University of Massachusetts, Amherst[2]

**Department of Computer Science
University of Colorado, Boulder[3]

†TRW
Redondo Beach, California

---

[1]This work was supported in part by National Science Foundation grants CCR-8451421, CCR-8704311, and CCR-8521398, TRW, and Hughes Aircraft.

[2]This work was supported in part by the following grants: Rome Air Development Center, F30602-86-C00006; National Science Foundation, DCR-8404217, DCR-8408143, DCR-8318776, and CER DCR-8500332; Control Data Corporation, No.84M103.

[3]Work supported by National Science Foundation grants MCS-8302644 and DCR-8403341, the Department of Energy grant #1537612, and the American Telephone and Teleglraph Company.

Abstract

Early software environments have supported a narrow range of activities (*programming* environments) or else been restricted to a single "hard-wired" software development process. The Arcadia consortium's joint research is focused on the construction of software environments that are highly integrated, yet flexible and extensible enough to support experimentation with alternative software processes. This has led us to view an environment as being composed of two distinct, cooperating parts. One is a *variant* part, consisting of a process definition plus the tools and objects specific to the particular defined process. The other is a fixed part, or *infrastructure*, providing an invariant basis on which various process definitions, and their associated tools and objects, can be supported. The major components of the infrastructure are a process programming language and interpreter, user interface facilities, and an object management subsystem. The process programming facility allows precise definition and automated support of software development and maintenance activities. The object management subsystem provides persistent storage for highly structured, typed objects, and supports the type system of the process programming language. The user interface management system makes the system hospitable for human users, who carry out the creative parts of process programs.

Keywords: environment, software process, process programming, object management, user interface management system

Contents

1 A Characterization of Software Environments 2

2 Software Process Definition and Interpretation 5
2.1 Software Processes 5
2.2 Process Programming Languages 7

3 Object Management 11

4 Interface with the Human User 18

5 Summary and Conclusion 23

References 25

1 A Characterization of Software Environments

A common thread that runs through the literature on software environments is that the purpose of environments is to ***support the user*** in some software development or maintenance activity. Sometimes this activity is highly constrained and well defined, such as constructing syntactically correct source code. Other environments have broader scope, but are highly restrictive in the order of events that are permitted. Still other environments are simply collections of tools and data management facilities that are believed to be helpful in a broad arena of software evolution activities.

Thoughtful consideration of this notion of "supporting the user" yields some important insights. First, if the activities that an environment supports are not precisely and unambiguously described, it is difficult for potential users to assess whether their needs will be met. Second, the facilities provided by an environment may be so loosely structured that they *could* support a variety of activities, but if all structuring and composition is solely the end user's responsibility, for which no automated support is provided, then undue burden is placed upon the user. It is likely that such an environment will be used to support only the simplest and smallest activities. Third, clearly specifying and supporting a specific activity may not be enough. Change to virtually any software development or maintenance activity is inevitable. Users will wish to incorporate new tools and development methodologies. If the environment's structure is closely entwined with the original activity, then accommodating change may be difficult and costly, or not possible at all.

In our estimation, therefore, a useful environment will

- support clearly and precisely defined activities,
- mechanize the structuring and composition of support functions, and
- accommodate changes and personal preferences.

Such notions of usefulness could be applied to any software system, however, and environments must be distinguished from large, multi-option tools if the term is to have any useful meaning. We believe that this distinction cannot be made for some early so-called programming environments. We therefore offer the following definitions of next generation software environments.

An environment consists of two parts: a fixed part and a variant part. The aspects of an environment that are prone to change, such as the tool set, belong to the variant part. The unchanging mechanisms necessary to ensure the integrity, extensibility, flexibility, and integration of the environment are encapsulated in the fixed part. We believe this division is a critical separation of concerns. In contrast, environments like *Interlisp* [53] and *Smalltalk* [20] make no distinction between fixed

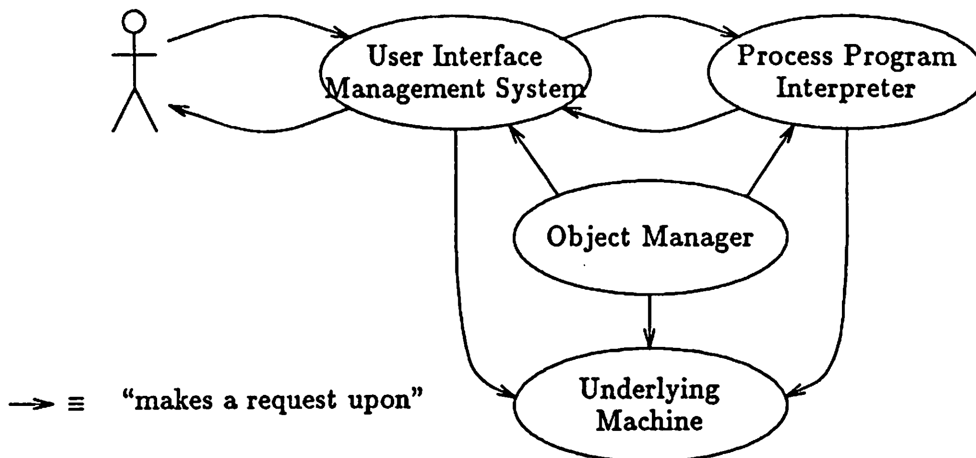

Figure 1: The active entities of an environment infrastructure with the "makes a request upon" relationship shown between them. The variant parts of an environment are not shown.

and variant part, or even between the environment and the software developed within the environment (to the point where the software can run nowhere except within the environment — they are inextricably bound).

More specifically, the variant part consists of the evolving set of data objects (such as specifications, programs, and test data) along with rigorous, detailed descriptions of software development or maintenance activities, which we term "process programs" [33] [34]. These activities are defined in terms of individual specific operators, which correspond to the classical notion of tools. These operators can themselves be modeled as process programs (if not actually implemented as such), and are correspondingly in the variant part too.

The fixed part, which can also be termed the *environment infrastructure*, consists of all the mechanisms necessary for the automated interpretation of process programs. Specifically, it consists of a language for writing process programs, an agent that enables interpretation of process programs, including mechanisms for managing persistent objects, and facilities for providing interfaces to the human user. The fixed part also encapsulates assumptions regarding its underlying machine. That is, it may make use of an operating system, a storage manager, and so forth. These assumptions are also components of the infrastructure. The components of the infrastructure that are active entities are shown in Figure 1.

The user may initiate action by making a request for support of some activity. The request is passed along, through the user interface management system, to the process program interpreter. Assuming that the request is well formed, the interpreter initiates a series of actions, executing the process specified by the user.

This execution will involve invoking operators upon operands — both of which are managed by the "object manager" component. Examples of operators include lexers, parsers, code generators, debuggers, test data generators, and specification and design language processors. Examples of operands include source code, executable modules, test data, specifications, designs, management data, symbol tables, various internal representations of programs, designs, system generation directives (e.g., make files), and intermediate analysis results.

Many of the operators will themselves be processes that cause additional actions to occur. Note that the entities belonging to the variant part of the environment, such as tools and data objects, are *managed* by this part of the infrastructure — they are called into action at appropriate times and do their work — but they are not themselves components of the infrastructure.

There are, of course, operations in a software process that can only be performed by people. In essence, the mundane aspects of processes are automated in this view of environments, while the creative aspects are performed by creative agents — people. Accordingly the interpreter may issue a request for an operation to be performed by the user, passing the request through the user interface management system.

The user interface management system may directly request services of the object manager for storage of windows and depictions of objects. Finally, the underlying machine provides services for each of the other three automated components of the infrastructure.

In short, the infrastructure is a virtual machine for the interpretation of process programs. (It is for this reason that the interpreter is part of the infrastructure, rather than a tool belonging to the variant part.)

Clearly we are burying many of the critical and interesting technical issues inside the process interpretation system and the object manager. The subsequent sections of this paper, which are organized around the entities shown in Figure 1, will clarify and elucidate the ideas suggested here. The notions of software processes, process programming languages, and the interpretation of process programs are considered first, in Section 2, as they are central to our view of environments. These are followed by discussions of object management in Section 3 and the user interface management system in Section 4. Additional details of these issues, plus discussion of the necessary capabilities of the underlying machine, measurement and evaluation of environments, and development and technology transfer can be found in [50].

2 Software Process Definition and Interpretation

2.1 Software Processes

Perhaps the most striking feature of the environment architecture described here is that it empowers users to rigorously specify their software products and product types *and* rigorously and explicitly specify alterable process programs to guide in the development and maintenance of these products. Previous environment architectures have exploited only primitive notions of explicitly specified products and processes. They have supported relatively fixed processes and products, often specified only implicitly. Moreover, the user's freedom to specify the process supported or the type of product produced by the environment was generally sharply restricted.

For instance, some previous environments have been aimed at supporting the development and maintenance of software specifications and designs. Systems such as *PAISLey* [62], *RSL/REVS* [3], *SARA* [18], Data Flow Diagram Designs [58], Jackson System Development [10], and *USE* [59], are examples of such previous environments. While certainly valuable for their intended purposes, each of these systems provides support for the creation of a relatively narrow range of software objects by relatively restrictive and inflexible processes. Specifically, they guide users to the development of design or specification objects in a particular fixed discipline and format, which is usually pictorial or graphical. For example, *RSL/REVS* is organized to strongly aid users in creating, analyzing, and maintaining designs as hierarchies of graph structures that are heavily annotated. In such environments, the exact structure of the objects and their pictorial representations vary from one system to another. In some cases the user is able to tailor and adapt these software object types. Invariably, however, these adaptations can be made only within a narrow range. For example, users of such environments may be able to select the specific fields to be incorporated into a design node, but only from a given fixed list of fields and types.

In addition, these design and specification support environments attempt to lead the user through specific procedural processes that are intended to expedite the creation of designs and improve the chance that the resulting designs are well formed and in compliance with the guidelines of the design or specification methodology being supported. Accordingly, such environments are often either indifferent or overtly hostile towards attempts to create design objects of new or different types, or to follow development procedures that have been devised by the user. From our perspective, these environments contain "hard coded" object specifications and processes (which they effectively support). They are not hospitable to user attempts to make significant alterations to such processes or design object specifications.

There are also other environments whose goals are aimed more at supporting the development of code. Environments such as *Interlisp* [53], *Arcturus* [45], and *Cedar* [52] integrate facilities to support the creation of code in specific languages. They

support such common activities as editing, parsing, debugging, and documentation. These environments assume that user activities can be uniformly and smoothly integrated by viewing them as examination and transformation of a single uniform representation of one product — code or a representation of the code. In Interlisp all software products, as well as the procedures and tools used to create them, are considered to be instances of lists. In Arcturus, software products and the commands used to manipulate them are all instances of Ada code. As long as the user's activities are effectively modeled in these ways, these environments provide strong support. As the user seeks to model software products and processes as objects of different types, support from these environments falters and becomes awkward.

Similarly, intelligent editors such as the *Cornell Program Synthesizer* [51], *Integral-C* [40], *Gandalf* [22], and *Mentor* [17] are all effective in integrating user activities, but only over a restricted range. Here, the integration rationale is that user activities all revolve around a parsed representation of code in a specific language. Experience has shown that this representation supports many user activities more effectively than a textual representation of the code. Shifting focus from text to an internal representation, however, does not solve the problems posed by restricting users from being able to create and manipulate software objects of types of their own creation using explicit processes of their own creation. Structure editors implicitly assume that users are concerned with a few types of objects.

In all of these cases the effectiveness of the support tools is drastically reduced when the process that the user wishes to carry out is not anticipated by the environment. In the case of code synthesis using environments such as *Mentor* or *Interlisp*, when the user attempts to execute process steps operating on object types not related to code, such as tests, support is weak. In the case of a design environment, when the user attempts to stray beyond the supported methodology, or attempts to carry out such processes as coding or testing, support similarly is weak. The objective here is not to criticize specific systems, but to point out that the value of any *environment* is closely related to its abilities to support all the user's activities.

Support for only limited, pre-determined processes is particularly disturbing in view of the observation that there is currently little consensus about what constitutes adequate software products and effective software processes, and that products and processes must therefore be expected to vary from user to user, from location to location, and from time to time. The most effective process architecture for a spread-sheet application, for example, will be different from the most effective process architecture for developing a complex command-control-communications system. Similarly, project schedule, budget, personnel, reliability, or portability constraints will strongly condition the most effective choice and sequencing of major process activities. Just as the programming of a software product is more effective when preceded by product requirements analysis, architecture definition, and design

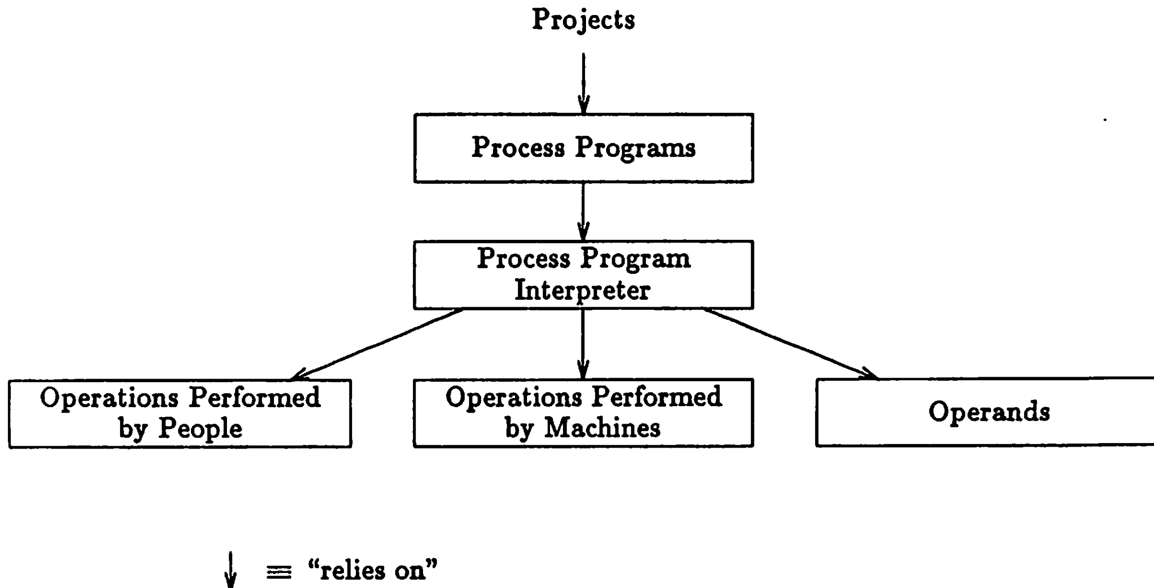

Figure 2: The relationship of process programs to the projects they support and to the mechanisms that implement them.

activities, so will be the programming of one's software lifecycle process.

Many observers believe that progress towards understanding what constitutes adequate products and effective processes can only follow from experimentation with alternatives. We believe that the best way to facilitate such experimentation is to enable users to easily yet rigorously describe software products and processes in ways that are convenient and effective and to support the rapid interpretation of processes in terms of software tools and procedures. From our perspective it is clear that this is tantamount to the creation of environments in which the variant part — i.e., the product specifications, process descriptions and set of operators — is specifiable by the user, and in which the environment exploits this specification to fashion user support, utilizing components from the fixed part.

2.2 Process Programming Languages

In this section we discuss some key details of how process programs should be used to enable users to flexibly specify how they wish to have machine resources applied to the support of their activities. Figure 2 shows the relationship of process programs to the projects they support and the operations and operands that implement them. We see that projects are to be directly supported by specifications of how they are to be carried out, where these specifications are to be captured by process programs. Process programs rely upon a process program interpreter to carry out

their commands. This interpreter is responsible for translating the specifications embodied in the process program into operations on operands. As indicated in Figure 2, some of the operators upon which the interpreter relies are executed by computing devices, but others are executed by humans. An important aspect of process programming is that it is a vehicle for indicating which activities are to be carried out by humans and how these activities are to be coordinated with activities carried out by computing devices. We believe that one of the most important objectives of software engineering is to orchestrate the way in which humans, support software systems, and machines are to be coordinated to isolate and specify problems, to attack their solution and to determine the degree to which these activities have been successful or need to be modified. These activities are neither completely mechanical and automatable, nor completely spontaneous and indefinable. Rather they must be a careful blend of these approaches. We believe that this blend can best be specified and communicated by expressing it in a concrete form — namely the process program.

Early Precedents and Lessons All software projects have as their goal the creation and/or modification of software products. They work towards this goal by carrying out software operations on software operands. Thus, at the most basic level, process programs must be viewed as vehicles for specifying the coordination of such operations. It is worthwhile to observe that even operating system control languages can be viewed as primitive process programming languages. Language processors and system facilities are legitimate operators and the files managed by the operating system's file system are the operands. The user issues commands to the operating system and it effects the requested operations. Thus, command files or scripts are primitive process programs, using the operating system command language as the process programming language.

It is important to observe that these primitive process programs are used to indicate the ways in which human operations are to be coordinated with software operations. For example, users employ operators incorporated into tools, such as browsers, to help them carry out (human) selection operations. Selected objects are then used as operands to subsequent software operators, such as edit and compile. As another example, users often carry out standard sequences of operations at certain fixed times during software projects. They may invoke scripts to automatically compile new code, or automatically check consistency of new code with support libraries. These scripts orchestrate the interaction between machine operations (e.g., compiling and checking) and human operations (e.g., creating code). In this way, the scripts are small but good examples of process programs.

Scripts have also been used to automatically create new objects and maintain certain types of consistency between new and existing objects. For example, scripts are used to automatically recompile source code when support libraries have been

changed, or to recreate executables when source code has been changed. The *Make* system in *Unix* [19] is an example of a capability whose goal is to facilitate the creation of powerful scripts of just this sort, through the use of a terse and precise notation. Clearly this notation is a process programming language, albeit limited.

Current Issues and Needs Operating system command languages and *Make* can be used to write process programs, but they lack the power needed to effectively program large and complex processes. One of their most basic and significant deficiencies is their weak facilities for defining software objects as instances of types. Our ideas about the need for an object typing facility, and the way in which it should be provided, are elaborated more fully in the next section. Suffice it to say here that typing offers a powerful vehicle for organizing not just the basic objects to be managed in a software project, but also for defining and organizing the operators — both human and machine executable — to be employed by the project[1].

Going further, we believe the next basic capability that a process programming language must support is specification of the order in which operators are to be applied to operands. Many operating system command languages incorporate some flow of control operators, but these are usually quite primitive, often consisting only of basic looping and alternation constructs (*csh* is an exception here). In fact it seems that paucity of control flow expressive power may well be the weakest aspect of most operating system command languages.

Interestingly, other early attempts at rigorously expressing software process have focused directly on these aspects. Most notable among these attempts have been efforts to use diagrammatic representations to depict the major features of large-scale software processes. In this work, principal software processes were represented by boxes, and flow of control relations among them were represented by arrows. The "Waterfall Model" of software development relied upon this device in an early attempt to represent an overall software development process [41]. Almost immediately, software process modelers attempted to use these pictorial representations to also show other relations such as data flow or process hierarchy. Even later work attempted to simultaneously show diagrammatic representations of many key relations among a variety of types of software objects. In order to do this, data objects were differentiated from process objects by making distinctions between the shapes of the boxes representing them. Distinctions among relations were made by defining different shapes of arrowheads and different colors and shapes of arrows to represent these different relations. Some examples of such more advanced diagrammatic process representations are ETVX boxes [37], SADT diagrams [39] and Software Development Graphs [5].

[1] In either case, the semantics of operations can be formally defined using, for example, pre- and post- conditions. These conditions, in turn, utilize the accessing primitives that participate in defining the object types.

Inevitably these efforts are limited by the fact that there are arbitrarily many valid relations among the large number of software objects required to adequately describe software processes, and that different users may at different times wish to study various combinations of them. Creating one single diagram containing all of these relations is hardly a solution, as such a diagram is so complicated as to confound all understanding. Creation of a single internal representation capturing all of the complex relations in a software process, and then relying upon tools to draw specific diagrams ("views") upon request, seems to be a plausible solution to this dilemma. We believe that a process program, written in a suitable language, is the appropriate device for representing this.

Thus, we see that software practitioners have used operating system command languages as primitive programming languages to program micro-level processes. These languages are too primitive to support effective expression of large and complex software processes. At the same time, modelers have attempted to portray these large and complex processes with diagram systems that have been unable to clearly and *precisely* express all of the facets and details of true software processes. They have been thwarted in their attempts for the same reason — namely the lack of a language that is suitably expressive. Thus two diverse and important currents both point towards the same basic problem — the need to define a process programming language.

Characteristics of a Process Programming Language We have already noted that software process programming requires the ability to define a wide variety of software object types, and that this is best supported by powerful data typing and relationship mechanisms. (The issues here are addressed in depth in Section 3.) Moreover, support for controlling the procedural flow inherent in software processes must also be provided. Our early attempts to construct process programs for realistic software processes have convinced us that the range of control flow operators required is quite broad. Clearly iteration, selection, and procedure invocation are required in order to accurately portray the way software processes are carried out. In addition such control flow capabilities as parallel execution and exception handling seem essential.

Further consideration of how to enable specification of flow of control raises the deeper issue of whether an imperative model is the most appropriate linguistic paradigm to use in process programming. Although many aspects of many kinds of software process seem to be inherently procedural and algorithmic, there are other software activities that defy simple algorithmic description and suggest that the declarative paradigm is much more appropriate. Design creation is an example of such an activity.

In design creation the goal is to create a design specification. Often (e.g., in the case of the Software Cost Reduction methodology [35]), it is quite possible to

specify the goal object — namely a complex structure of carefully prescribed design elements — but it is not clear how to give complete procedural details on how to construct it. In such cases it is often reasonable to create rules that guide and constrain activities, such as the selection of good candidates for design elaboration, or that can intelligently raise issues about apparent inconsistencies among design elements. Thus some aspects of design seem to be rule-based. Other aspects, such as the orderly elaboration of details of design elements and their correlation with each other, are much more procedural. This suggests that a process programming language might ideally be a language that combines procedural and rule-base paradigms.

Furthermore, our early work indicates that an important aspect of software processes is that they often create other processes that are executed later on. For example, test planning is a process whose goal is the creation of another process that is to be executed at some future point in the execution of the software development process. During test planning, the test plan is created as a software object. This may entail such subactivities as development of test cases, encoding of algorithmic strategies for the systematic execution of the test cases, and development of procedures for capturing test results. Much later in the development process, after code has been developed, this test plan object must be "executed." This entails treating the test plan object as a process, rather than an operand. This passive/active nature of some software processes points to the desirability of a language in which code and data are freely interchangeable, as in Lisp.

In the Arcadia project we are experimenting with software process programming languages. In our earliest efforts we are coding process modules in a variety of language paradigms, attempting to arrive at a more precise set of requirements for this language. (In the interests of uncluttered exposition, this section has assumed that there is only a single process programming language. Different process programming languages (not application languages) would be associated with distinct environments. We recognize, however, that development of a single language that is appropriate for describing all software processes may be technically infeasible. Even if it is feasible, achieving consensus may be impossible. In any event, environments and process programming languages must evolve over the years. Until these issues are resolved we must accommodate multiple process languages in our research endeavor. Our simplistic assumption used in the presentation must therefore yield to a more robust view in practice. The two issues of evolution and research methodology are discussed very briefly in the concluding section.)

3 Object Management

An environment user's primary objective is to create and/or maintain a *software product*. No matter what process program might be used in creating and maintaining

it, a software product will be a very complex and highly interrelated collection of objects. Those objects will be of widely different kinds, ranging from source code and executable modules to documentation and test plans. Each kind of object will have an associated set of applicable operations, but operations applicable to one kind of object will generally not be appropriate for use with other kinds. This suggests that an environment's fixed part should provide support for managing typed objects and a rich set of relationships among them.

As Figure 1 indicates, the object management system will be a major component of the Arcadia environment infrastructure. It will be responsible for managing objects in two distinct classes: the *components* of the software products being produced by users of the environment, and the *tools and information structures* that constitute the environment itself. From the process programming perspective, the former can be viewed as the (input and output) data manipulated by a process program while the latter are the operators and internal data structures of the process program. (As previously noted, an object can move back and forth between categories during its lifetime.)

The object management system will provide the underlying mechanism on which the data management capabilities of a process programming language and a process program interpreter can be constructed. A particular process programming language might present its users exactly the same object management capabilities that the environment's object management system provides, as an assembly language presents its users exactly the same data types provided by the underlying machine. It seems equally likely, however, that a process programming language might offer a different view of objects than that provided by the environment's object manager. In either case, the properties of the object management system will influence the data management aspects of an environment's process programming languages.

Most environment builders have had to rely on a traditional file or database system for storing the objects associated with their environment. It is our belief, however, that a much richer set of capabilities for controlling object creation, access, and organization is essential to an environment. In particular, a suitably powerful object management system will enhance the environment's support for change, its integration, its support for software reuse, and its support for cooperative work by multiple developers.

Work on environments during the last decade has revealed four important areas of concern that must be addressed by an object management system. These are:

- type systems;
- relationship sytems;
- object persistence; and

- concurrent and distributed object management.

Each poses interesting problems. The capabilities sought in each of these areas and the problems we foresee are discussed below.

Type systems As indicated above, we view a type system as the primary mechanism for describing and maintaining objects. The object manager should be able to enforce the type system, hiding the internal structure of typed objects behind well-defined interfaces and strictly controlling the operations that can be performed on those objects. If all objects are instances of abstract data types, it is easier to share objects or to change their implementations. Thus, basing the object management system on a typing system that fully supports data abstraction will contribute to environment flexibility and software reuse.

Current approaches to object management in environments fall far short of providing full support for typed objects. Typically the components of a product are treated simply as files [19] and tools are viewed as operators applicable to the contents of those files. Usually in such systems, only a predetermined and limited number of different kinds of components (e.g., source file, object file) and operations (e.g., compiler, linker) are available. The *Odin* subsystem of *Toolpack* [15] improved on this simple view by using file name extensions as a weak form of typing mechanism for files (going beyond the capabilities of *Make*). It also allowed users to define which tools could operate on or produce files of various types. The *System Modeller*, developed as part of the *Cedar* system [26] used the term "object" for referring to the files containing product components, but did not treat the objects as instances of abstract types. The Common APSE Interface Set (CAIS) [9] defines a system model with three kinds of nodes — file, structural, and process — but does not treat those nodes as typed objects. While clearly improving on the simple use of files, all of these systems provide only partial support for typed objects. Meanwhile, work on support for typed objects within the traditional database community [48,11,63], while encouraging, is still in its primitive stages and far from providing the flexibility and power needed for object management in a software development environment [4]. Recent work on rich type systems, particularly in the context of object-oriented languages [29], is also encouraging, but also still in its infancy. No consensus has yet emerged on a desirable and appropriate set of features for such a type system.

Thus, one major object management research issue is: What kind of type system is needed to describe the objects populating a software development environment? The type system needs to be flexible and powerful enough to capture the relevant properties of environment objects. Tools, processes, and perhaps even types themselves need to be treated as typed objects. Once the capabilities of the type system are clearly delineated, suitable mechanisms for realizing those capabilities must be found. While there are many intriguing proposals for type mechanisms, it is not

clear which of these (e.g., single vs. multiple inheritance, specification vs. representation inheritance, generics, static vs. dynamic binding) form a compatible set providing the capabilities needed for environments.

Relationship systems Closely related to the ability to precisely define and maintain the typed objects in the environment is the ability to capture and maintain the relationships among those objects. Much environment work in the last ten years has focused on mechanisms for describing, reasoning about, or exploiting relationships among objects. Examples of relationships include those connecting various versions of a module, or those between the modules constituting a configuration, or those between a module and all the others that it calls, or those joining activities in a work breakdown structure. Examples of tools that reason about or exploit relationships among objects include revision control systems [55], automated system building tools [19], call graph analyzers, and work activity management systems [21].

Explicitly indicating the relationships among an environment's tools and information structures should make it easier to modify the environment since the effect of changes can be determined. Moreover, capabilities that rely on relationships, such as inference and derivation, will enhance environment integration by allowing users to interact with the environment at a high level, leaving the intermediate steps to be automatically determined. Generic relationship capabilities will also enhance integration by providing a uniform set of capabilities across different kinds of relationships.

A weakness in previous work is that there has been no systematic treatment of the numerous and complex relationships that exist among environment objects. The CAIS notions of primary and secondary relationships (also found in the node structure of the *ALS* [54]), *Odin*'s derivation graphs, the PMDB work of Penedo and Stuckle [36], and the system models of *Cedar* represent important starting points. The concept of configuration threads found in *DSEE* [27] and the relationship capabilities for module interconnection languages provided by *INTERCOL* [56] are additional examples of partial treatments specialized to one class of relationships.

Thus, another important object management research question is: What are suitable primitives and constructors for defining the relationships needed in environments? It is not clear whether the diverse relationships needed in software development can be captured in a single model or not. Moreover, how should the relationship structure and the type system interact? Associated with the relationship system is a set of capabilities, such as consistency checking, derivation tracking, and inferencing. Work needs to be done on identifying these capabilities and in exploring how generic such capabilities can be. For example, can generic consistency checking tools applicable to the relationship structure subsume the specialized consistency analyses associated with interface control or configuration management?

Another important concern is when and how such capabilities are initiated. Some must be requested by the environment user, either directly or via an executing process. Others can be more effective if triggered by resulting events. Thus, support for "active" objects or daemons that are triggered by process or user-specified events in the environment is needed.

Object persistence The object manager must be able to preserve the components of software products and the constituents of software environments for arbitrary periods of time. Moreover, it should preserve both the structure and the restrictions on how these objects can be manipulated that are imposed by the type system. Under such a scheme, the traditional distinction between primary and secondary storage representations of objects is hidden within the typed object abstraction. This can free both environment users and environment builders from concern about distinctions between internal and external representations of objects and conversions between those representations. Thus, the object manager should support *persistence*, enabling objects to continue to exist beyond the lifetime of any of the tools that manipulate them and preserving the integrity of their types and relationships to other objects.

Current approaches to persistence, based on files or databases, require explicit action by the tools. Using a file system, a tool must take responsibility for converting the internal form of an object to an acceptable (e.g., linear) external form and, when needed, converting it back. There are few restrictions to assure that the type of an object is not violated (e.g., that its contents are not altered using an editor while it resides in the file) or changed (e.g., that a stack is not read back as an array). Using a database system, the tool must make calls on the database to explicitly store and retrieve information. Current databases provide support for only a limited number of types, so once again the tool must provide the conversion algorithms and there is no guarantee of type integrity. There has been some interesting work on merging database support into programming languages [2,16,32], although implemented prototypes have been very restrictive about the supported types [2] or the underlying program model [16].

Thus, an important issue that must be addressed by the object manager is: How should persistence be provided for arbitrarily complex, typed objects? To permit maximum flexibility in the creation of objects and their relationships, the persistence of an object should be a property orthogonal to all other object properties. It is not clear how persistence should be recognized in a program (e.g., declared as part of the type or explicitly requested with the instantiation of an object) or how invisible persistence can be (e.g., no need to explicitly "commit" or "linearize" objects). Supporting a rich type system and providing an invisible line between memory and secondary storage raise challenging problems.

Concurrent and distributed object management To allow multiple users to work conveniently on the same software development project requires support for concurrent and distributed object management. Assuming a network of workstations, different members of a development project may simultaneously be invoking the same or different tools to operate on one or more of the same objects. Thus, the object manager must be able to mediate concurrent use of objects and to maintain consistency of both the objects and their relationships. Ideally, the object manager should make the distributed nature of the object base and the concurrent access to its objects invisible to users and tools in the environment.

A variety of approaches for handling distribution and concurrency have emerged from programming language [1] [24] [28] and file system and database research [23] [57] [42]. Unfortunately no single model for dealing with these issues is universally accepted within one of these domains, let alone for objects that move between them. Moreover, some of the more popular approaches are ill-suited for use in an environment object management system. Locking schemes, for example, typically apply to entire objects and do not permit concurrent access to disjoint subsets of an object's components, which may be a frequent occurrence in an environment. Transactions schemes generally presuppose relatively short duration transactions, while a software developer's transactions (e.g., update a source program) may last for days or weeks. The rollback approach to conflict resolution is also of questionable value in an environment.

Thus, one of the major problems facing object management is: What are appropriate constructs for expressing distribution and concurrency constraints and what underlying mechanisms must be provided to support these constraints? It is not clear what storage management primitives need to be provided to adequately capture the distribution and concurrency needs of an environment. As with types, relationships, and support for persistence, the appropriate descriptive notations must be developed as well. Also, where should the desired concurrent/distributed behavior be described — in the tools that create the actual instances of the objects, in the abstract data types that define the objects, or in the process programs that describe how the objects are to co-exist within the environment?

Arcadia Approaches As indicated above, much work has previously been done on problems related to object management. That work, however, has generally been directed toward solving individual problems, leading in some cases to incompatible solutions, and has not yet resulted in consensus on the appropriateness of those solutions. Moreover, much of the work has been oriented toward domains other than software development.

The approach to developing an object management system that is being taken in the Arcadia project is therefore one of synthesis and extension. In particular, we are initially looking to programming language technology for guidance in the design

of a type system and the expression of distribution and concurrency constraints, and initially looking to database technology for guidance in the design of mechanisms for persistence, relationships, change, and distribution.

It is clear that some new solutions are still required to satisfy the special needs of software development environments. To sharpen our understanding of these needs, we are examining process programs for a wide variety of activities, examining a wide variety of tools that would make use of the object management system, and reflecting on our experience building *Odin* and *Keystone* [14], which can be viewed as primitive object management systems. We are also developing formal models, as we have previously done for module interface relationships [60], for describing and evaluating the various capabilities intended for inclusion into the object management system.

One design that we are considering provides a functional layering of the desired capabilities. At the lowest level are facilities for such things as storage management, concurrency control, and transaction management. The next level supports the basic concept of types, essentially defining the type system provided by the object manager. Above that are primitives for realizing object relationships. The capabilities for revision control, partitioning of objects into libraries, and the like, appear at the highest level. All of this together provides the basis on which type systems for process programming languages can be implemented.

We intend to build successively more sophisticated prototypes of the object management system. This activity will be facilitated by the recent trend in database research toward the development of *database toolkits* [43,4,44,11]. These toolkits provide basic, low-level capabilities such as storage management, concurrency control, and transaction management. The idea is to provide a foundation upon which to build higher-level capabilities, such as those for typing and relationships. The obvious benefit of using such a foundation is the ability to experiment with alternative higher-level structures without having to construct instances of the lower-level facilities for each such alternative. These toolkits are intended to be "general purpose" and we intend to experiment with prototypes of the toolkits to ensure that our particular needs can be satisfied.

Until database toolkits become available, we are building prototypes that examine particular aspects of object management. Two examples of this are a relationship management system and an application generator called *Graphite* [13]. The relationship management system is intended as a vehicle for exploring the suitability of various automated constraint-satisfaction and inferencing techniques in the domain of process programming. In particular, it provides a general framework for specifying goals in terms of relationships over objects, and mechanisms, such as backward and forward chaining, for reasoning about the satisfaction of those goals. *Graphite* is being used to investigate issues in the specification of types, insulating tools from changes to those types, and hiding details of how instances of those types are made

persistent. The class of types that *Graphite* focuses on is attributed graphs, since it is clear that this particular class is important in a software development environment. For instance, one of our uses for attributed graphs is to internally represent programs.

4 Interface with the Human User

The user interface management system is the third major component of an environment's fixed part. We consider it here, discussing first the characteristics of good interfaces that an environment should exhibit. Some outstanding problems are then noted. The remainder of the section addresses various specific approaches to the design of user interface management systems, including separating tool functionality from interaction properties, the use of abstract depictions in managing the display, and techniques to aid in achieving uniformity.

Characteristics of good interfaces Broad consensus exists on the qualities which distingish good user interfaces for software environments. *Uniformity* (or *consistency*) reduces the difficulty of learning new activities and moving between activities. The *direct manipulation* interaction paradigm, using graphics and pointing devices, increases the communication bandwidth between tool and user. *Permissive* (or *non-preemptive*) interfaces allow the user to interleave activities in a natural way.

Uniformity reduces the number of details a human user must remember, and increases skill transfer between activities. A uniform interface makes the same set of operations available everywhere they make sense, and allows the user to specify an operation in the same manner wherever it is available. Interpreter-based programming environments made significant early progress toward uniformity by unifying the command language and programming language of the environment. More recently, editor-based programming environments have provided a uniform set of commands for manipulating program source code, blurring the distinction between editing, compiling, and debugging. Limited progress has been made in providing a uniform interface across a wider variety of activities, mostly by imposing informal standards (like the *Macintosh* user interface guidelines [25]) and providing a toolkit of reusable components (scrollbars, menus, and the like).

Direct manipulation, or more precisely the illusion of directly manipulating a set of objects, requires a rich visual representation of state. This visual representation unburdens the users' short-term memory, replacing recall tasks with easier recognition tasks. (Menus serve a similar purpose with respect to remembering commands). Objects are referred to with a pointing device and through implicit pointing (e.g., cursor position.) Changes in the representation provide immediate confirmation of user actions. The basic principles of direct manipulation are applicable to character

displays, but modern bitmapped workstations are capable of richer visual representations of state. Pioneering work in the application of graphics to programming and software engineering include the *Incense* debugging system [31], the *Balsa* algorithm animation system [8], and the *Pecan* programming environment [38].

Permissiveness is an essential aspect of direct manipulation, too seldom achieved in current systems. A permissive interface allows the user to choose the next action, arbitrarily interleaving interactions with each object depicted on the screen. The converse of permissiveness is *preemption.* A preemptive interface imposes an order on user actions. The prompt/input paradigm of gathering input is a classic example of preemption.

Window systems are primarily a means of limiting preemption. Windows grafted onto a conventional system in the form of multiple virtual terminals provide a minimal degree of permissiveness, sufficient for the user to temporarily escape from the control of a single application. The multiple views of *Pecan* [38] and the *Pi* debugger [12] hint at the richer interaction possible when each tool may coordinate several threads of control.

Outstanding problems Uniformity becomes both more important and harder to maintain as the scope of an environment grows. A large, extensible environment will contain tools contributed by a diverse community of developers and users. Both the toolset and interaction techniques can be expected to evolve during the lifetime of the environment. A critical problem, then, is decoupling the human interface from tools so that each may evolve independently. Providing a set of reusable components is helpful, but may not be enough by itself. The *SunView* facilities [49], for instance, encourage similar visual appearance across tools, but they are not much help in establishing consistent interpretations of mouse and keyboard actions within windows managed by tools. The interface between interaction and tool functionality (in the application domain) is the most troublesome interface in modularizing interactive graphics programs. Because graphics toolkits deal entirely with the graphical domain, they do not help clean up this interface.

The problem becomes apparent when one notes that other tools, as well as human users, may use a tool component. A good human interface is generally not a good tool interface. An all-purpose interface, like Unix character streams, is unlikely to be satisfactory in either role. Thus, in current Unix-based systems, the set of tool-usable tools is quite disjoint from the set of interactive tools. It is difficult, for instance, to use a screen-oriented editor or a spreadsheet program as part of a pipe or shell script.

Techniques and approaches User interface management systems (UIMS) is an active area of research, outside the context of software environments research as well as within it. The following paragraphs discuss current approaches to separat-

ing application functionality from interaction facilities, managing the display, and establishing a uniform interface to all the functions supported by an environment. The design of the *Chiron* user interface subsystem of *Arcadia* is briefly presented as an example of a system that brings together several threads from current user interface research.

Separating functionality and interaction. Several current approaches to direct manipulation interfaces carefully separate the application domain (or *model*) from the presentation domain (or *view*). A tool manipulates objects in an application domain. An encapsulated tool component (sometimes called the *controller*) maintains consistency between objects in the two domains, so that the presentation domain accurately reflects the state of application objects and the application domain properly responds to user activity in the presentation domain. We see this as a key separation.

In "editor" environments supporting a narrow set of objects and functions, a tool component may map the central application data structure (typically a parse tree) to the presentation domain. Separation of concerns between application domain and presentation domain is achieved, but at the cost of requiring all environment facilities to operate on the shared data structure rather than a variety of data structures suited to different applications. Environments of wide scope require a more flexible scheme.

The *Incense* debugger [31] introduced the notion of *artists* to maintain the depictions of each particular type of application data. Each artist encapsulates information about a particular application data type, and how it should be represented. Since this information is encapsulated in individual artists, outside of any shared user interface infrastructure, tools are not forced to share a common representation or data model for application objects. This is important, but raises the question of associating artists with types.

Multiple inheritance in a type system provides a powerful mechanism for associating artists with application data types. The *annotation* mechanism of *Loops* [47] [46] can be used to trigger an artist whenever an application object is manipulated. Lisp object systems [7,30,6] provide a similar capability with method-mixing in multiple inheritance. Conceptually, an artist is "wrapped around" an existing data type, as illustrated in Figure 3, so that the old interface (available operations) shows through the new.

Managing the display. Current approaches to user interfaces generally interpose an intermediate level of representation between application objects and their concrete depiction on the screen (Figure 4). This *abstract depiction* serves several purposes. It is generally more convenient for an artist to manipulate a structured description of a display than a lower-level representation, especially if the structure

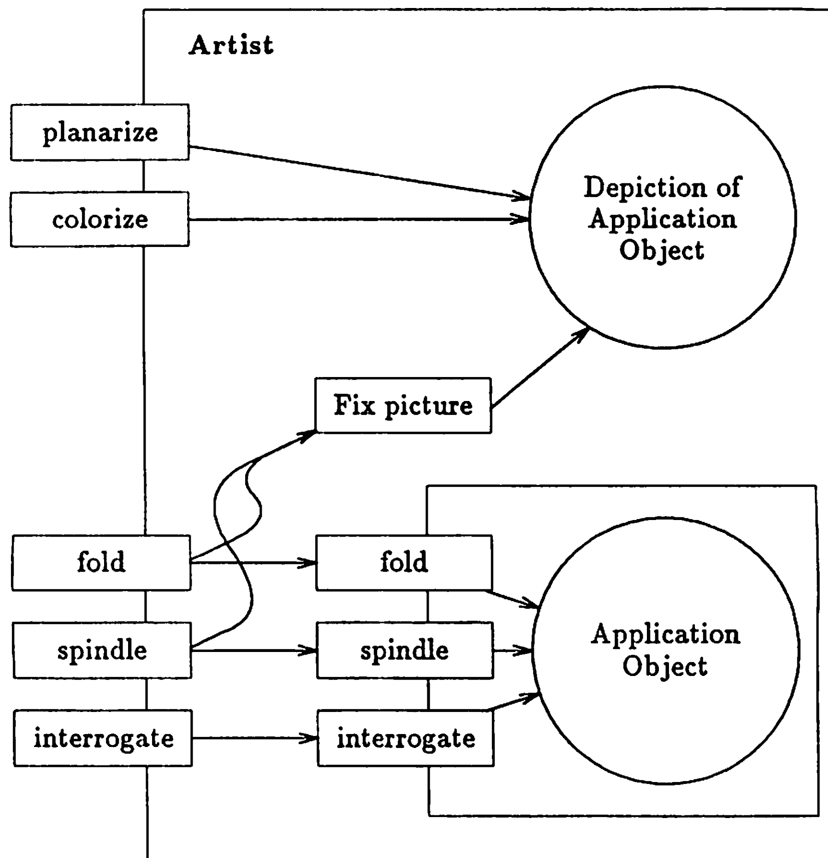

Figure 3: An artist is logically "wrapped around" an abstract data type. The application (or *model*) object is encapsulated in an abstract data type, with visible operations `fold`, `spindle`, and `interrogate`. Each of these operations "show through" the artist, in the sense that the artist exports operations with identical signatures and semantics, except that the presentation object (or *view*) is updated as a side effect. Operations which do not change the application object (`interrogate`, in this diagram) are simply re-exported without change; additional operations on the presentation object (`planarize`, `colorize`) may be added.

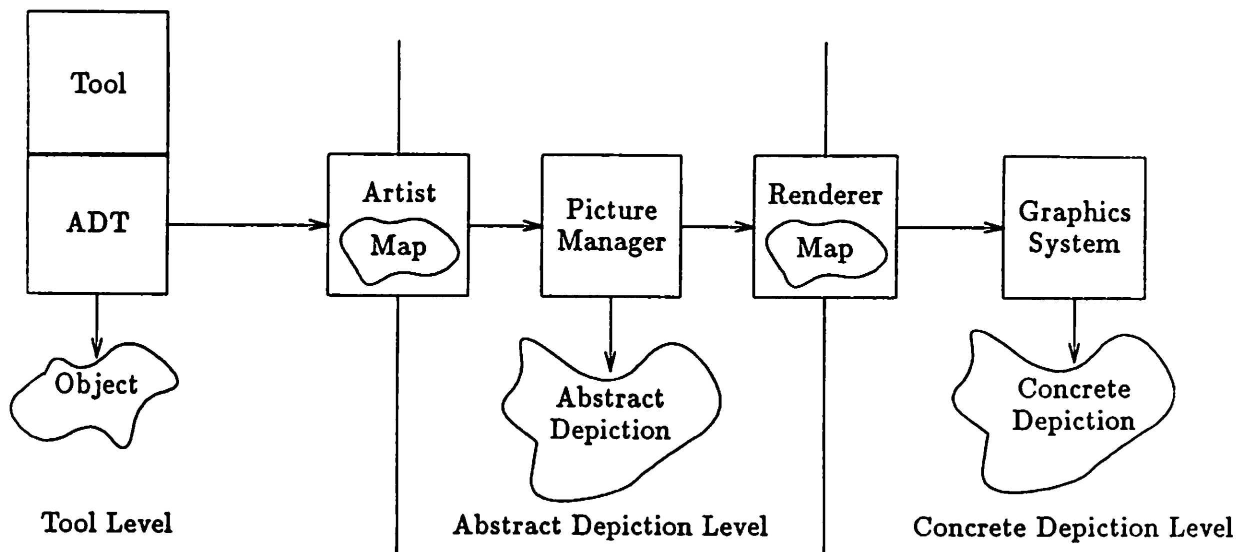

Figure 4: The most difficult part of interactive graphics programming is maintaining the association between application objects and their depictions. When a structured intermediate representation (an *abstract depiction*) is interposed between application objects and their depictions, this task can be considerably simplified. The system can maintain the relationship between the abstract and concrete depictions, and associate input events with particular components of the abstract depiction. The artist that created that particular component can then be notified; it need only maintain the relationship between high-level graphical objects and the application objects they depict.

of the abstract depiction reflects the structure of the depicted object. Also, manipulating a portion of the abstract depiction can result in efficient incremental update of the display, provided the rendering agent is able to determine which portions of the concrete depiction may be affected by the change.

More importantly, an abstract depiction can be used as a basis for input correlation, relating an input action (e.g., mouse click) with a particular application object. Window systems provide input correlation down to the window level, but not within windows. This is sufficient for menus and scrollbars, which can be designed so that each choice lies in its own window, but not sufficient for general diagrammatic depictions of software objects. An abstract depiction level managed by the environment can perform input correlation to the level of individual picture elements.

Approaches to uniformity. Centralized interpretation of low-level input can be used to achieve a basic level of uniformity. For instance, if the lexeme *select* is bound to a single click of the leftmost mouse button, then the application will receive the event *select*, rather than a raw key click, when the button is pressed. Binding of lexemes to raw events should always be under control of the user, rather than the tool builder. Techniques adequate for administering this level of interpretation are well known (e.g., the TIP tables of *Cedar*). Central administration can also guarantee consistent interpretation of a small set of "global" commands, for instance, terminating a tool. Anyone who has attempted to kill an unfamiliar Unix program with keyboard incantations will appreciate the importance of such guarantees.

Reusable components are a complementary approach to promoting uniformity. Application-independent components, such as scrollbars, are already in common use. Clean encapsulation of interaction facilities makes it feasible to provide reusable components for data abstractions in a particular application domain (e.g., Petri nets), as well. Since artists are associated with abstract data types, the path of least resistance for tool developers is to reuse an artist for all interactive tools dealing with a particular data abstraction.

Arcadia approach to user interface. The *Chiron* user interface subsystem of *Arcadia* is characterized by artists bound to abstract data types through a type inheritance mechanism, a simple diagram-oriented abstract depiction, concurrency between and within tools, and support for uniformity across tools.

Since abstract data types are key to modularizing tool fragments in *Arcadia*, *Chiron* uses type inheritance to bind artists to objects. An artist inherits application functionality and adds new state (a depiction) and new operations (e.g., `planarize`, `colorize`). It also manages side effects to the new state from existing operations (i.e., updating the display when the object changes). *Chiron* provides a diagram-oriented $2\frac{1}{2}$D hierarchical display model, including nested and overlapping windows. Artists manipulate this *abstract depiction*. *Chiron* maps it into the *concrete depiction*, typically a bitmap.

Chiron emphasizes concurrency between and within tools. Each depicted object may have its own thread of control, and each may independently maintain its depiction and react to user actions. In addition, a rendering agent maintains the concrete depiction concurrently with manipulations of the abstract depiction (subject to interlocks on the latter), and input proceeds concurrently with output. Additional detail on *Chiron* can be found in [61].

5 Summary and Conclusion

The current flurry of activity in environments and in software process specification is exciting. A proper focus for environments — supporting the user's multiple,

complex activities — is being reemphasized at a time when some pertinent subtechnologies are maturing. This paper has presented definitions that are useful in categorizing and assessing developments in environments, and has attempted to separate some key concerns. In so doing, a number of emerging principles and important open problems have been identified and some promising research directions described.

One key distinction is between an environment's fixed infrastructure and its variant part. As part of the infrastructure, a user interface management system provides communication between humans and executing software processes. These processes are described in a formal process programming language and are interpreted by a process program interpreter. Mundane, automatable activities are handled directly; creative activities are performed by creative agents: people. A key component of the automated interpreter is an object management subsystem, whose typing system, relationship system, persistence scheme, and facilities for distributed and concurrent object management, support the constructs of the process programming language. Having process programming as a key part of the concept makes the environment an active agent, rather than a purely reactive one.

As noted earlier, for purposes of exposition this paper has promoted a rather static view of the process programming language and has identified an environment as being associated with only one process programming language. It is clear, however, that with experience we will want to revise the language definition and we may want to support multiple process languages. Such change is a concern and must be both anticipated and supported. In the same way that we want environments to support processes that manage an evolving collection of software products, and to support the evolution of processes themselves (ideally even as they are executing), we also want to support the evolution of process programming languages. Better understanding of the forcing functions here could lead us to revise our separation between fixed and variant parts. We find it rewarding to note that description of an environment architecture that supports the first two kinds of change exposes the next meta-question, namely change in the process programming language itself.

In our estimation, progress on the various fronts of environment research is now tied to realistic prototype development, empirical evaluation, and technology transfer. Prototypes are needed to validate concepts, generate feedback, and provide demonstrations that new environment technologies are useful to large development teams tackling large development activities. To be fully convincing, and to generate as much insight as possible, realistic prototypes must be subjected to well-designed empirical evaluation. Carefully planned technology transfer activities are then needed to ensure that the sought-after benefits are fully realized.

The Arcadia consortium has been formed to do research in environment architectures. We are attempting to make major strides in the development of the fundamental technologies, develop prototypes, conduct careful empirical studies, and

move the technology to industrial practice. Joint funding to support the activities of the consortium began in August 1987.

Acknowledgements

Arcadia has benefited from the ongoing encouragement of William L. Scherlis and Stephen L. Squires.

We gratefully acknowledge the contributions of participants from each institution of the consortium, who have been instrumental in shaping virtually every aspect of Arcadia. During the past two years key contributions have come from D. Baker, B. Boehm, A. Brindle, D. Fisher, D. Heimbigner, C. LeDoux, M. Penedo, I. Shy, S. Sykes, W. Tracz, and S. Zeil.

Additional contributions have been made by R. Adrion, G. Barbanis, M. Burdick, G. Clemm, R. Cowan, E. Epp, S. Gamalel-Din, C. Kelly, S. Krane, R. King, D. Luckham, T. Nguyen, D. Richardson, W. Rosenblatt, R. Schmalz, C. Snider, S. Sutton, P. Tarr, and D. Troup.

References

[1] *Military Standard Ada Programming Language (ANSI/MIL-STD-1815A-1983)*. American National Standards Institute, January 1983.

[2] M. Atkinson, P. Bailey, K. Chisholm, P. Cockshott, and R. Morrison. An approach to persistent programming. *The Computer Journal*, 26(4):360–365, 1983.

[3] T. E. Bell, D. C. Bixler, and M. E. Dyer. An extendable approach to computer-aided software requirements engineering. *IEEE Transactions on Software Engineering*, SE-3(1):49–60, February 1977.

[4] P. A. Bernstein. Database system support for software engineering. In *Proceedings of the Ninth International Conference on Software Engineering*, pages 166–178, IEEE Computer Society Press, Monterey, California, March 1987.

[5] D. Bjorner. On the use of formal methods in software development. In *Proceedings of the Ninth International Conference on Software Engineering*, pages 17–29, IEEE Computer Society Press, Monterey, California, March 1987.

[6] D. G. Bobrow, L. G. DeMichiel, R. P. Gabriel, K. Kahn, S. E. Keene, G. Kiczales, L. Masinter, D. A. Moon, M. Stefik, and D. L. Weinreb. Common Lisp object system specification. April 1987. Draft document, X3J13 Committee for the Common Lisp Standards effort.

[7] D. G. Bobrow, K. Kahn, G. Kiczales, L. Masinter, M. Stefik, and F. Zdybel. CommonLoops: merging Lisp and object-oriented programming. In *OOPSLA '86: Object Oriented Programming Systems, Languages, and Applications*, pages 17–29, 1986. Appeared as *SIGPLAN Notices 21*(11).

[8] M. H. Brown and R. Sedgewick. A system for algorithm animation. *Computer Graphics*, 18(3):177–186, July 1984.

[9] *Military Standard Common APSE Interface Set (CAIS), Proposed MIL–STD–CAIS*. Department of Defense, January 1985.

[10] J. R. Cameron. An overview of JSD. *IEEE Transactions on Software Engineering*, SE-12(2):222–240, February 1986.

[11] M. J. Carey, D. J. DeWitt, D. Frank, G. Graefe, J. E. Richardson, E. J. Shekita, and M. Muralikrishna. *The Architecture of the EXODUS Extensible DMBS: A Preliminary Report.* Technical Report CS-644, Computer Science Department, University of Wisconsin – Madison, Madison, May 1986.

[12] T. A. Cargill. The feel of pi. In *Winter 1986 USENIX Technical Conference*, pages 62–71, USENIX Association, Denver, Colorado, January 1986.

[13] L. A. Clarke, J. C. Wileden, and A. L. Wolf. Graphite: a meta-tool for Ada environment development. In *Proceedings of the IEEE Computer Society Second International Conference on Ada Applications and Environments*, pages 81–90, IEEE Computer Society Press, Miami Beach, Florida, April 1986.

[14] G. M. Clemm, D. Heimbigner, L. Osterweil, and L. G. Williams. Keystone: a federated software environment. In *Proceedings of the Workshop on Software Engineering Environments for Programming-in-the-Large*, pages 80–88, June 1985.

[15] G. M. Clemm and L. J. Osterweil. *A Mechanism for Environment Integration.* Technical Report CU-CS-323-86, University of Colorado, Boulder, Jan. 1986.

[16] *CLF Overview.* USC Information Sciences Institute, Marina del Rey, California, March 1986. Informal Report.

[17] V. Donzeau-Gouge, G. Huet, G. Kahn, and B. Lang. Programming environments based on structured editors: the Mentor experience. In *Interactive Programming Environments*, pages 128–140, McGraw-Hill Book Co., New York, 1984.

[18] G. Estrin, R. S. Fenchel, R. R. Rasouk, and M. K. Vernon. SARA (System ARchitects Apprentice): modeling, analysis, and simulation support for design of concurrent systems. *IEEE Transactions on Software Engineering*, SE-12(2):293–311, February 1986.

[19] S. I. Feldman. Make—a program for maintaining computer programs. *Software — Practice & Experience*, 9(4):255–265, Apr. 1979.

[20] A. Goldberg and D. Robson. *Smalltalk-80: The Language And Its Implementation.* Addison Wesley, 1983.

[21] C. Green, D. Luckham, R. Balzer, T. Cheatham, , and C. Rich. *Report on a Knowledge-Based Software Assistant.* Technical Report, Kestrel Institute, June 1983.

[22] A. N. Habermann and D. Notkin. Gandalf: software development environments. *IEEE Transactions on Software Engineering*, SE-12(12):1117–1127, Dec. 1986.

[23] D. Heimbigner and D. McLeod. A federated architecture for information management. *ACM Transactions on Office Information Systems*, 3(3):253–278, July 1985.

[24] C. A. R. Hoare. Communicating sequential processes. *Communications of the ACM*, 21(8):666–677, 1978.

[25] *Inside Macintosh.* Apple Computer, Inc., Cupertino, California, promotional edition, March 1985.

[26] B. W. Lampson and E. E. Schmidt. Organizing software in a distributed environment. In *Proceedings of the ACM SIGPLAN '83 Symposium on Programming Language Issues in Software Systems*, pages 1–13, June 1983.

[27] D. B. Leblang and J. Robert P. Chase. Computer-aided software engineering in a distributed workstation environment. *SIGPLAN Notices*, 19(5):104–112, May 1984. (Proceedings of the First ACM SIGSOFT/SIGPLAN Symposium on Practical Software Development Environments).

[28] B. Liskov and R. Scheifler. Guardians and actions: linguistic support for robust, distributed programs. *ACM Transactions on Programming Languages and Systems*, 5(3):381–404, July 1983.

[29] N. Meyrowitz, editor. *Object-Oriented Programming Systems, Languages, and Applications Conference Proceedings*, ACM SIGPLAN, Portland, Oregon, September 1986. (Appeared as SIGPLAN Notices, 21(11), November 1986.).

[30] D. A. Moon. Object-oriented programming with flavors. In *OOPSLA '86: Object Oriented Programming Systems, Languages, and Applications*, pages 1–8, 1986. Appeared as *SIGPLAN Notices 21*(11).

[31] B. A. Myers. Incense: a system for displaying data structures. *Computer Graphics*, 17(3):115–125, July 1983.

[32] J. A. Orenstein, S. K. Sarin, and U. Dayal. *Managing Persistent Objects in Ada*. Technical Report CCA-86-03, Computer Corporation of America, Cambridge, Massachusetts, May 1986.

[33] L. Osterweil. *A Process-Object Centered View of Software Environment Architecture*. Technical Report CU-CS-332-86, University of Colorado, Boulder, May 1986.

[34] L. Osterweil. Software processes are software too. In *Proceedings of the Ninth International Conference on Software Engineering*, pages 2–13, Monterey, California, March 1987.

[35] D. L. Parnas and P. C. Clements. A rational design process: how and why to fake it. *IEEE Transactions on Software Engineering*, SE-12(2):251–257, February 1986.

[36] M. H. Penedo and E. D. Stuckle. PMDB — A project master database for software engineering environments. In *Proceedings of the Eighth International Conference on Software Engineering*, pages 150–157, London, Aug. 1985.

[37] R. A. Radice, N. K. Roth, A. C. O. Jr., and W. A. Ciarfella. A programming process architecture. *IBM Systems Journal*, 24(2):79–90, 1985.

[38] S. P. Reiss. PECAN: program development systems that support multiple views. *IEEE Transactions on Software Engineering*, SE-11(3):276–285, 1985.

[39] D. T. Ross and K. E. S. Jr. Structured analysis for requirements definition. *IEEE Transactions on Software Engineering*, SE-3(1):6–15, February 1977.

[40] G. Ross. Integral C — a practical environment for c programming. In *Proceedings of the Second ACM SIGSOFT/SIGPLAN Symposium on Practical Software Development Environments*, pages 42–48, December 1986.

[41] W. W. Royce. Managing the development of large software systems. In *Proceedings, IEEE WESCON*, pages 1–9, IEEE, August 1970. also reprinted in Proceedings 9th International Conference on Software Engineering, pp. 328–338.

[42] M. Satyanarayanan, J. H. Howard, D. A. Nichols, R. N. Sidebotham, A. Z. Spector, and M. J. West. The ITC distributed file system: principles and design. In *Proceedings of the Tenth ACM Symposium on Operating Systems Principles*, pages 35–50, 1985.

[43] A. H. Skarra, S. B. Zdonik, and S. P. Reiss. An object server for an object-oriented database system. In *Proceedings of the Object-Oriented Database Systems Workshop*, pages 196–204, IEEE Computer Society Press, Sep. 1986.

[44] A. Z. Spector. *Distributed Transaction Processing in the Camelot System*. Technical Report CMU-CS-87-100, Computer Science Department, Carnegie-Mellon University, Pittsburgh, Pennsylvania, January 1987.

[45] T. A. Standish and R. N. Taylor. Arcturus: a prototype advanced Ada programming environment. *Software Engineering Notes*, 9(3):57–64, May 1984. (Proceedings of the ACM SIGSOFT/SIGPLAN Software Engineering Symposium on Practical Software Development Environments).

[46] M. Stefik and D. Bobrow. Object-oriented programming: themes and variations. *AI Magazine*, 6(4):40–62, Winter 1986.

[47] M. J. Stefik, D. G. Bobrow, and K. M. Kahn. Integrating access-oriented programming into a multiparadigm environment. *IEEE Software*, 3(1):10–18, Jan. 1986.

[48] M. Stonebraker and L. A. Rowe. The design of POSTGRES. In *Proceedings of the ACM SIGMOD '86 International Conference on Management of Data*, pages 340–355, June 1986.

[49] *SunView Programmer's Guide*. Sun Microsystems, Inc., Mountain View, California, February 1986.

[50] R. N. Taylor, D. A. Baker, F. C. Belz, B. W. Boehm, L. Clarke, D. A. Fisher, L. Osterweil, R. W. Selby, J. C. Wileden, A. L. Wolf, and M. Young. *Next Generation Software Environments: Principles, Problems, and Research Directions*. Technical Report 87–16, Department of Information and Computer Science, University of California, Irvine, July 1987. Also issued as Arcadia Document Number UCI–87–10, University of Colorado Technical Report Number CU–CS–370–87, University of Massachusetts, Amherst Technical Report Number 87–63, and Incremental Systems Corporation Technical Report Number 87–7–1.

[51] T. Teitelbaum and T. Reps. The Cornell Pprogram Synthesizer: A syntax directed programming environment. *Communications of the ACM*, 24(9):563–573, Sep. 1981.

[52] W. Teitelman. A tour through Cedar. *IEEE Transactions on Software Engineering*, SE-11(3):285–302, 1985.

[53] W. Teitelman and L. Masinter. The Interlisp programming environment. *Computer*, 14(4):25–33, April 1981.

[54] R. M. Thall. The KAPSE for the Ada Language System. In *Proceedings of the AdaTEC Conference on Ada*, pages 31–47, ACM, October 1982.

[55] W. F. Tichy. Design, implementation, and evaluation of a revision control system. In *Proceedings of the Sixth International Conference on Software Engineering*, pages 58–67, Tokyo, Japan, Sep. 1982.

[56] W. F. Tichy. Software development control based on module interconnection. In *Proceedings of the Fourth International Conference on Software Engineering*, pages 29–41, IEEE Computer Society Press, Munich, West Germany, Sep. 1979.

[57] B. Walker, G. Popek, R. English, C. Kline, and G. Thiel. The LOCUS distributed operating system. In *Proceedings of the Ninth ACM Symposium on Operating Systems Principles*, pages 49–70, 1983.

[58] P. T. Ward. The transformation schema: an extension of the data flow diagram to represent control and timing. *IEEE Transactions on Software Engineering*, SE-12(2):198–210, February 1986.

[59] A. I. Wasserman, P. A. Pircher, D. T. Shewmake, and M. L. Kersten. Developing interactive information systems with the user software engineering methodology. *IEEE Transactions on Software Engineering*, SE-12(2):326–345, February 1986.

[60] A. L. Wolf, L. A. Clarke, and J. C. Wileden. A formal model for describing and evaluating visibility control mechanisms. In *Proceedings of the IEEE Computer Society 1986 International Conference on Computer Languages*, pages 182–189, IEEE Computer Society Press, Miami Beach, Florida, October 1986.

[61] M. Young, R. N. Taylor, and D. B. Troup. *Software Environment Architectures and User Interface Facilities.* Arcadia document Arcadia Document UCI-87-07, Department of Information and Computer Science, University of California, Irvine, May 1987. Submitted for publication.

[62] P. Zave and W. Schell. Salient features of an executable specification language and its environment. *IEEE Transactions on Software Engineering*, SE–12(2):312–325, February 1986.

[63] S. B. Zdonik and P. Wegner. Language and methodology for object-oriented database environments. In *Proceedings of the Nineteenth Annual Hawaii International Conference on System Sciences*, pages 378–387, Jan. 1986.