

Distribution and Abstract Types in Emerald

A. Black
N. Hutchinson
E. Jul
H. Levy
L. Carter



Distribution and Abstract Types in Emerald

ANDREW BLACK, NORMAN HUTCHINSON, ERIC JUL, HENRY LEVY, AND LARRY CARTER

Abstract—Emerald is an object-based language for programming distributed subsystems and applications. Its novel features include 1) a single object model that is used both for programming in the small and in the large, 2) support for abstract types, and 3) an explicit notion of object location and mobility. This paper outlines the goals of Emerald, relates Emerald to previous work, and describes its type system and distribution support. We are currently constructing a prototype implementation of Emerald.

Index Terms—Abstract data types, distributed operating system, distributed programming, object-oriented programming, process migration, type checking.

I. INTRODUCTION

WHILE distributed systems are now commonplace, the programming of distributed applications is still somewhat of a black art. We believe that the complexity of distributed applications is heightened by the lack of programming language support for distribution. For example, most distributed applications are implemented by calling operating system communications primitives, such as send and receive. The programmer is responsible for locating the communications target, explicitly packaging parameters, and so on. Before the introduction of concurrent programming languages, concurrent programs were constructed in a similar fashion. Language support for concurrency greatly simplified concurrent programming; we believe that language support for distribution can have a similar effect on distributed programming. Experience with the remote procedure call facilities of Cedar/Mesa [2] and with the Eden Programming Language [1] has justified this belief. With Emerald, we intend to go beyond simple syntactic support for message send and receive, and address some of the fundamental semantic problems of distribution.

Although distribution has many benefits [22], it also introduces challenges for the designer of a distributed language. First, the language must present a model of distributed computation; it must provide the conceptual framework that allows the programmer to define the objects that he manipulates in both the local and distributed

environment. Second, it must provide for both intra- and internode communication in an efficient manner. The semantics of communication and computation should be consistent in the local and remote cases. Third, it must allow the programmer to exploit the inherent parallelism and availability of a distributed system. Fourth, since shutting down and recompiling an entire distributed system in order to modify some component is unacceptable, the language must permit system extensibility without recompilation; existing programs must continue to work in collaboration with new programs.

Our research focuses on simplifying the programming of distributed subsystems and applications by providing language support for distribution. We have designed an object-based language, called *Emerald*, and a distributed run-time system for Emerald that facilitate the construction of distributed programs for a local area network of independent nodes (workstations). The novel features of Emerald include: 1) a single object model that is used for both programming in the small and in the large, 2) support for abstract types, and 3) an explicit notion of object location and mobility. The goal of our research is to demonstrate the feasibility of using one simple semantic model for programming both sequential, single-node applications, and concurrent, potentially distributed applications. We currently have a prototype Emerald compiler and run-time system running on a local area network of VAX® workstations.

The next sections present a discussion of previous work in distributed programming languages and an overview of Emerald. Following sections describe the type system and the support for distribution.

II. REVIEW OF PREVIOUS SYSTEMS

To date, languages have supported distribution in several different ways. In the Xerox Cedar System [33], a remote procedure call facility allows programs to access remote servers through standard Cedar/Mesa procedure calls [2]. The advantage of this approach is that it requires no change to the semantics of the language. Automatically generated stub routines on the client and server machines are responsible for packing and unpacking parameters and transmitting and receiving messages. Programmers access a remote service in the same way that they would access a local service, except that they must explicitly locate and connect to the service before it can be used.

®VAX is a registered trademark of Digital Equipment Corporation.

Manuscript received January 31, 1986; revised June 16, 1986. This work was supported in part by the National Science Foundation under Grants MCS-8004111 and DCR-8420945, by the University of Copenhagen, Denmark, under Grant J.nr. 574-2,2, and by a Digital Equipment Corporation External Research Grant.

A. Black, N. Hutchinson, E. Jul, and H. Levy are with the Department of Computer Science, University of Washington, Seattle, WA 98195.

L. Carter is with the IBM Thomas J. Watson Research Center, Yorktown Heights, NY 10598. This work was performed while he was a visitor at the University of Washington.

IEEE Log Number 8611363.

TABLE I
EMERALD LANGUAGE FEATURES

	Emerald	Argus		Xerox RPC	Eden		Smalltalk
		CLU	Guardians		EPL	Objects	
Uniform Object Model	✓	✓					✓
Distribution	✓		✓	✓		✓	
Mobile Objects	✓					✓	
Direct Objects	✓	✓		✓	✓		✓
Concurrency Control	✓		✓	✓	✓	✓	
Type Checking	✓	✓	✓	✓	✓		
Abstract Types	✓			✓			

At the University of Washington, the Eden Programming Language (EPL) [3] has been developed for writing distributed applications on the Eden system [1]. EPL is an extension of Concurrent Euclid [19] that provides location-independent invocation of Eden's objects in an integrated distributed system. Because the location of an Eden object is (conceptually) evaluated at each invocation, objects are free to move at any time; one need not locate or connect to an object before invoking it.

The M.I.T. Argus system [24], [25] is an ambitious distributed language project that extends the CLU language [23] to support atomic transactions in a distributed environment. An Argus *Guardian* encapsulates the notion of a physical machine. Inside a guardian are data objects and processes. One Guardian communicates with another by calling a handler in the target Guardian, to which data are passed by value.

In general, object-based systems and languages have viewed their objects in two ways: as large, long-lived resources (e.g., files) as in operating systems such as Hydra [35] and StarOS [20], or as small resources (e.g., records and integers), as in languages such as CLU [23] and Smalltalk [15]. In a distributed environment, both views seem to have their place; the Argus and EPL languages each support two kinds of objects. Argus has Guardians, which are network-wide objects, and CLU objects, which are local to a Guardian; EPL has network-wide Eden objects that contain local EPL variables, monitors, and modules. The reason for this dichotomy is one of locality and performance; local objects communicate through shared store, while network objects communicate through message passing, which requires more communications overhead. Unfortunately, this requires the programmer to use two different object abstraction mechanisms, to code in two different styles, and to foresee all possible uses to which an object will be put. For example, while programming a Collaborative Editing System in Argus, Greif *et al.* [16] have observed that a designer can be forced to use a Guardian where a cluster might be more appropriate.

Emerald has drawn on the experience of all of these systems. The most important difference between Emerald and these systems is Emerald's uniform model of computation. Like Smalltalk, all entities in Emerald are ob-

jects, and a single semantic model suffices to define them. Unlike Smalltalk, however, Emerald is a distributed programming language; its object model is sufficient to describe both local data objects and potentially remote objects containing independent processes. Table I enumerates the principal features of Emerald and compares them to those of Argus, Xerox RPC, EPL, and Smalltalk. The following section provides a brief overview of the Emerald programming language, focusing on its uniform object model, and following sections deal with two important aspects of Emerald: its type system which is based on the concept of abstract types, and its support for distribution.

III. INTRODUCTION TO EMERALD

Emerald is object-based and all information is encapsulated in objects. An object model is appropriate for a distributed system because it implicitly defines 1) the units of distribution and movement, and 2) the entities that communicate. All objects in Emerald are coded using the same object definition mechanism, regardless of the way in which they will be used. The Emerald object model is appropriate for defining *small* objects such as integers, characters, and Booleans as well as *large* objects such as directories and compilers. While different objects may be represented by the system in different ways, all objects exhibit the same semantics. Each Emerald object exports a set of operations; an object can be manipulated only by invocation of one of those operations. Furthermore, Emerald objects are *mobile*. Objects can move at any time, and can be invoked without knowledge of their location.

Each Emerald object has four components:

- 1) An *identity*, which distinguishes the object from all others within the network.
- 2) A *representation*, which consists of the data stored in the object. The representation of a programmer-defined object is composed of a collection of references to other objects.
- 3) A set of *operations*, which defines the functions and procedures that the object can execute. Some operations are exported and may be invoked by other objects, while others may be private to the object.
- 4) An optional *process*, which operates in parallel with

```

const inCoreFile == object inCoreFile
  export Read, Seek, Write
  monitor
4   const maximumSize == 200
   const CharacterVector == Vector.of[Character]
   var contents : CharacterVector
   var position : Integer

8   operation Read → [c : Character]
   if position <= contents.upperbound then
     c ← contents.getElement[position]
     position ← position + 1
12  else
     c ← nil
   end if
  end Read

16  operation Seek[p : Integer]
   position ← p
  end Seek

20  operation Write[c : Character]
   assert position <= contents.upperbound
   contents.setElement[position, c]
   position ← position + 1
24  end Write

  initially
   contents ← CharacterVector.create[maximumSize]
   position ← 0
28  end initially
  end monitor
end inCoreFile

```

Fig. 1. An inCoreFile object.

invocations of the object's operations. An object with a process has an active existence and executes independently of other objects. An object that has no process is a passive data object and executes only as a result of invocations.

An Emerald object also has several attributes. An object has a *location* that specifies the node on which that object is currently resident. Emerald objects can be defined to be *immutable*. Immutability is an assertion on the part of the programmer that the abstract state of an object does not change; it is not a concrete property and the system does not attempt to check it. On the other hand, the system takes advantage of immutability by copying such objects on remote reference.

Fig. 1 shows an Emerald definition of a simplified file object. This object supports the usual read, seek, and write operations expected from files, but of only a single character at a time. Its representation consists of a vector of characters and a current position indicator. Its three operations, *Read*, *Seek*, and *Write*, are exported and therefore available to users of the object. Since inCoreFile has no process, it is a passive data object and executes only as a result of invocations.

To exploit the inherent parallelism of distributed systems, Emerald supports concurrency both between objects and within an object. Separate threads of control are provided in the form of processes. Each object may have a *process section* specifying a parameterless, anonymous operation to be invoked asynchronously when the object has been initialized. Processes on the same processor execute in quasi-parallel with respect to each other and con-

currently with respect to processes located on other processors.

While an object has a single independent process, at any point in time multiple processes can be executing within a single object. This results from multiple invocations of an object's operations by other processes. Operations and variables may optionally be specified in a *monitor section* of the object. Processes executing monitored operations have exclusive access to the monitored variables and may synchronize using conventional *condition* variables [18]. An object's process normally executes outside the monitored section, but, like any other process, it can invoke the monitored operations should it need access to shared variables.

Each object has an optional *initially section*—a parameterless operation that executes exactly once when the object is created and is used to initialize the object's state. When the initially operation is complete, the object's process is started and invocations can be accepted.

Although Emerald has a single, uniform model of objects, objects are implemented in several different ways. An important goal of Emerald is to provide a very efficient implementation for objects. In order to accomplish this goal, the Emerald compiler chooses for each object an implementation style appropriate for its use. Some support of this type is available in CLU [31]. The Smalltalk language also has two implementation styles: one for primitive objects such as integers and arrays, and one for user-defined objects. While primitive object operations are relatively efficient in Smalltalk, there is no mechanism whereby users can define equally efficient types.

There are three styles of implementation available for Emerald objects. Standard types such as integer are usually implemented by a single word of storage and compiler-generated in-line operations. For example, the integer *add* operation is reduced to a single machine instruction. Objects that are local to another object can often be implemented by a compiler-allocated data area and their operations implemented as normal (Pascal-like) procedure calls, thus avoiding the more general invocation mechanism provided by the run-time kernel. Finally, objects that can move or be remotely referenced are implemented by a kernel-allocated data area and are indirectly referenced through an object table. Invocations of such objects are performed by a combination of compiled code and the run-time kernel. For example, if inCoreFile is used as a temporary file within another object it may be represented in storage with local pointers and manipulated by in-line code. If it is used to convey information between passes of a distributed compiler it will require a remote procedure call interface.

IV. TYPES IN EMERALD

Emerald is a statically typed language that supports abstract types. We first motivate our decision to make Emerald a typed language (Section IV-A) and our decision to support abstract types (Section IV-B). The *conformity* relation which forms the basis of the type system of Emer-

ald is then introduced, first informally in Section IV-C, and then rigorously in Section IV-D. Section IV-E compares our type system to that of other languages and provides reasons for the major differences, and Section IV-F discusses type system implementation considerations.

A. Why Types?

In standard *value-oriented* programming languages, the type system protects the user from misinterpretation of values, e.g., it prevents a character operation from being applied to an integer. In an object-oriented language the programmer is not allowed to apply operations to values at all; he may request only that an object perform an operation on *itself*. The point here is that the representation of the object is never handled directly by external operations; only the internal operations of the object are allowed access to its representation. As a consequence, neither Smalltalk objects, nor Eden objects, nor Emerald objects need be typed in order to guarantee their integrity. In contrast to Smalltalk, where identifiers do not have declared types, we have chosen to make Emerald a typed language, i.e., the programmer must associate a type with each identifier that she declares. We expect to see several advantages from this approach.

First is better detection and notification of programming errors. Many errors found in Smalltalk programming are of the "message not understood" variety. In Smalltalk this message can be generated only at run time, when an invocation is attempted on an object that does not implement it; it generally indicates that a type error has occurred. In other words, the object assigned to a variable is not of the required "type", because it does not support the required operation. In Emerald, we wish to detect such errors earlier. Because Emerald is typed, these errors will often be detected through static type checking at compile time rather than at run time. However, due to the flexibility demanded by our environment, we cannot do complete type checking at compile time, and some type errors will be detectable only at run time. Even when this is the case, type errors will be detected when a type-incorrect *assignment* is attempted rather than when invoking an unimplemented operation on an object. Also, rather than generating error messages that indicate that an operation was not implemented, we wish to generate messages that give a better indication of the reason for the operation not being understood. Making Emerald a typed language allows us to do this, generating messages that indicate that "object O is not of type T."

Second, we feel that type checking can reduce the cost of invocation. Since there is no possibility of "message not understood" errors at invocation time, it follows that no check is necessary. In addition, we have developed an alternative procedure for locating the code to execute in response to an invocation request (doing *method lookup* in Smalltalk terminology) that may provide better performance than existing schemes. This is made possible by the introduction of typing into the language and is described more fully in Section IV-F.

B. Abstract Types

Emerald was designed to be used in the construction of open systems, i.e., those where system-level objects may be created and added to a running system after the basic system is operational. This implies that the type system must be flexible enough to enable old code to invoke newly implemented objects, provided that these objects behave in the expected way. This flexibility and extensibility is provided by both Smalltalk and Eden, but neither of these systems are typed, as was noted above. Emerald wishes to retain this flexibility, but within the framework of a typed programming language.

In order to do this, Emerald objects are typed abstractly. An abstract type defines the interface of an object—the set of operations supported, their *signatures*, and (in principle) their semantics. An operation signature includes the operation name and the number, names, and abstract types of the arguments and results. Each object implementation defines a similar set of operations, but in addition provides 1) a concrete representation for the object and 2) code to implement each of the operations. Abstract types are themselves objects that export a *get-Signature* operation. For details on this aspect of the type system, see [4].

The relationship between abstract types and object implementations is many-to-one in both directions: each object may implement several abstract types, and each abstract type may be implemented by several different objects. Fig. 2 illustrates these relationships. The object *DiskFile* implements the abstract type *InputOutputFile*, the abstract types *InputFile* and *OutputFile* (which require only a subset of the *InputOutputFile* operations), and also the abstract type *Any* (which requires no operations at all). The abstract type *InputOutputFile* illustrates that an abstract type may have several implementations, perhaps tuned to different usage patterns. Temporary files may be implemented in primary memory (using *InCoreFile* objects) to provide fast access while giving up permanence in the face of crashes. On the other hand, permanent files implemented using *DiskFiles* would continue to exist across crashes.

This separation of specification and implementation provides the flexibility found in untyped languages such as Smalltalk and Eden, but within the framework of a strongly-typed language. We expect to take advantage of this flexibility in two ways.

- New objects satisfying old interfaces can be added at any time without changing existing objects. For example, consider a windowing system where *Window* is an abstract type. New windows with new functionality can be added, and the window manager can manipulate them without requiring any modifications, or even recompilation or relinking.

- An abstract type may have multiple implementations tailored to particular usage patterns, such as the file example of the previous section. In addition to multiple implementations explicitly provided by the programmer, the compiler itself may generate multiple implementations of

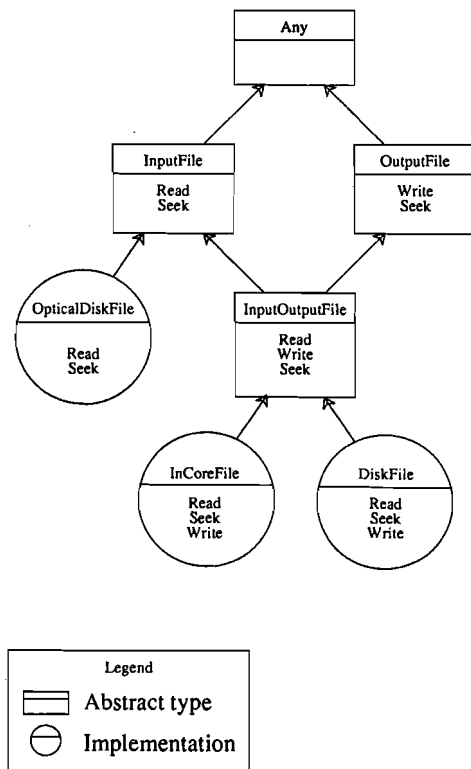


Fig. 2. Abstract types and object implementations.

some objects from the same source code. For example, one implementation might be appropriate for a remote object, another for the case in which the object and its invoker are known to be on the same machine, and yet another when the object is known to be private to some containing object.

An additional benefit of the separation of abstract types and implementations is that it is exactly the separation of specification and implementation necessary for top-down program design and separate compilation. In some existing languages, additional syntax and semantics are required to make possible this separation of specification and implementation. These include external declarations in Concurrent Euclid, Pascal, and Modula, and separate specification and implementation parts in Ada[®]. The type system of Emerald directly supports separation of specification and implementation; no additional language features are necessary.

C. The Emerald Type System

In Emerald, all identifiers are typed abstractly, i.e., the programmer declares the *abstract* type of the objects that an identifier may name. Such a declaration captures his knowledge of the set of invocations to which those objects should respond.

The notion of type *conformity* is central to Emerald. The legality of an assignment is based on the conformity of the assigned object and the abstract type declared for the identifier. This conformity will always be checked at

[®]Ada is a registered trademark of the U.S. Department of Defense (Ada Joint Program Office).

compile time, except where a run-time check is explicitly requested by the programmer. Roughly, a type P conforms to another type Q if P provides at least the operations of Q . (P may also provide additional operations.) Moreover, the types of the results of P 's operations must conform to the types of the results of the corresponding operations of Q . Finally, the types of the arguments of the corresponding operations must conform *in the opposite direction*, i.e., the arguments of Q 's operations must conform to those of P 's operations.

To illustrate the need for the parameter matching rules, consider the following examples. *Any* is the abstract type containing no operations, thus every type conforms to it.

```
const IntegerPusher ==
  type
    operation Push[Integer]
  end

const AnyPusher ==
  type
    operation Push[Any]
  end
```

These rather useless types define “write-only storage” into which integers and arbitrary objects can be *Pushed*. Intuitively, one would expect *AnyPusher* to conform to *IntegerPusher*, because an implementation of *AnyPusher* can be used wherever an *IntegerPusher* is required. The rules bear this out; the two types are identical except for the argument types of *Push*, and these conform in the opposite direction, i.e., *Integer* conforms to *Any*.

Now consider

```
const IntegerPopper ==
  type
    operation Pop → [Integer]
  end

const AnyPopper ==
  type
    operation Pop → [Any]
  end
```

Here *IntegerPopper* conforms to *AnyPopper*, because the results of *Pop* conform in the *same* direction. Finally, observe that

```
const IntegerStack ==
  type
    operation Push[Integer]
    operation Pop → [Integer]
  end
```

and

```
const AnyStack ==
  type
    operation Push[Any]
    operation Pop → [Any]
  end
```

are incomparable; they do not conform in either direction.

The reason for this should be obvious; users of an *IntegerStack* object expect its *Pop* operation to return an integer, so an *AnyStack* clearly will not do. Users of an *AnyStack* expect to be able to *Push* arbitrary objects; the *Push* of *IntegerStack* can be applied only to integers.

Emerald type conformity is similar to type compatibility in Owl [30], but differs in a fundamental way from inheritance in Smalltalk. In Smalltalk, a subclass does not necessarily conform to its superclass; for example, it may override some of the operations of the superclass so that they expect different classes of argument. Moreover, one class may conform to another without a subclass relationship existing between them. What a subclass and its superclass *do* have in common is part of their representation and some of their methods. In short, inheritance is a relationship between implementations, while conformity is a relationship between interfaces.

D. Formal Definition of Conformity

The above explanation of conformity in Emerald was not well-founded; the conformity of two types depended on the conformity of the types of the arguments and results of the operations defined by the types. This section will present a formal definition of the conformity relation between types, which forms the basis of Emerald's type system, as well as an algorithm for checking conformity. This presentation of conformity has been simplified in two ways. First, we have omitted operations which are functions and types which are immutable; these would only clutter the presentation with unnecessary detail. Second, we have ignored Emerald's support for polymorphism which is beyond the scope of this paper.

Let T be the set of type names, and let $a_i, r_j \in T$, for $0 \leq i \leq m, 0 \leq j \leq n$. A signature is either an expression of the form

$$a_1, \dots, a_m \rightarrow r_1, \dots, r_n$$

which has *arity* $\langle m, n \rangle$, or the distinguished *null signature*, Λ , on which *arity* is undefined. Let S be the set of all signatures, and F be the set of operation names. A *type declaration* of $t \in T$ is a total function (also denoted t) from F to S . Obviously, in actually declaring a type, one need only specify the operations which have non-null signatures. For example, the type *IntegerStack* can be defined as:

$$\text{IntegerStack}(x) \equiv \begin{cases} \text{Integer} \rightarrow & \text{if } x = \text{Push} \\ \rightarrow \text{Integer} & \text{if } x = \text{Pop} \\ \Lambda & \text{otherwise.} \end{cases}$$

Suppose Θ is a binary relation on T , i.e., $\Theta \subset T \times T$. Intuitively, Θ is a set of pairs which we hope are true assertions about conformity, i.e., $\langle t, u \rangle \in \Theta$ implies that t conforms to u . However, we have not defined conformity yet, so the intuition cannot be formalized. Now suppose that s, s' are elements of S . Then Θ induces a relation Σ on S by the following definition.

$$\langle s, s' \rangle \in \Sigma \text{ iff:}$$

i) $s' = \Lambda$, or

ii) $\text{arity}(s) = \text{arity}(s')$ and writing

$$s = a_1, \dots, a_m \rightarrow r_1, \dots, r_n$$

$$s' = a'_1, \dots, a'_m \rightarrow r'_1, \dots, r'_n$$

we have

$$\langle a'_i, a_i \rangle \in \Theta, \text{ for } i = 1, 2, \dots, m$$

$$\langle r_j, r'_j \rangle \in \Theta, \text{ for } j = 1, 2, \dots, n.$$

Informally, corresponding pairs of results must be in Θ and corresponding pairs of arguments must be in Θ^{-1} . Looking at the conformity of *IntegerPopper* and *AnyPopper* from the previous section, we have

$$s = \rightarrow \text{Integer}$$

and

$$s' = \rightarrow \text{Any}$$

In order for $\langle s, s' \rangle$ to be in Σ , $\langle \text{Integer}, \text{Any} \rangle$ must be in Θ .

Now, what is to distinguish an arbitrary Θ from our desired conformity relation? Exactly the requirement that the corresponding pairs of signatures are in Σ . Formally, we say that Θ is *valid* if, for all type names t and u , and all operation names $f \in F$

$$\langle t, u \rangle \in \Theta \Rightarrow \langle t(f), u(f) \rangle \in \Sigma.$$

Lemma: The union of two valid relations is valid.

Proof: Follows immediately from the definitions.

Corollary: There is a unique maximal valid relation, which we will denote by \triangleright .

Finally we have defined conformance. We write $\langle t, u \rangle \in \triangleright$ as $t \triangleright u$ or t conforms to u . One nice thing about valid relations all being contained in \triangleright is that we know that if two separately defined systems of declarations are each valid, then we will not get any surprises when we combine the declarations. Our concept of conformity is similar to the notion of implicit conversion of types as studied by Reynolds [28], [29].

Now we can define a *decision procedure* that will check whether a given statement of the form " $t \triangleright u$ " is true. Starting from the pair $\langle t, u \rangle$, we will build up the two relations Θ and Σ recursively. Θ will be a valid relation on T and Σ will be the relation on S induced by Θ . We will do this by, whenever we insert $\langle a, b \rangle$ into Θ , inserting pairs $\langle a(f), b(f) \rangle$ into Σ for all f such that $b(f) \neq \Lambda$. This ensures that Θ remains valid. Additionally, whenever we insert $\langle s, s' \rangle$ into Σ , we insert all the appropriate $\langle a'_i, a_i \rangle$'s and $\langle r_j, r'_j \rangle$'s into Θ so that Σ is indeed the derived relation for Θ . But we *fail* in attempting to insert a pair $\langle a(f), b(f) \rangle$ into Σ if the arities of $a(f)$ and $b(f)$ mismatch, or if $a(f) = \Lambda$ when $b(f) \neq \Lambda$. If we succeed, we will have constructed a valid relation. By the unique maximality of \triangleright , this will prove that $t \triangleright u$. On the other hand, we only inserted necessary elements into the relations Θ and Σ , so if the procedure *fails*, then t does not conform to u .

Let us apply this decision procedure to check whether $\text{IntegerStack} \triangleright \text{AnyStack}$. In order to insert $\langle \text{IntegerStack}, \text{AnyStack} \rangle$ into Θ , we must insert

$$\begin{aligned} & \langle \text{IntegerStack}(\text{Pop}), \text{AnyStack}(\text{Pop}) \rangle \\ & = \langle \rightarrow \text{Integer}, \rightarrow \text{Any} \rangle \end{aligned}$$

and

$$\begin{aligned} & \langle \text{IntegerStack}(\text{Push}), \text{AnyStack}(\text{Push}) \rangle \\ & = \langle \text{Integer} \rightarrow, \text{Any} \rightarrow \rangle \end{aligned}$$

into Σ . Inserting $\langle \rightarrow \text{Integer}, \rightarrow \text{Any} \rangle$ into Σ causes us to insert $\langle \text{Integer}, \text{Any} \rangle$ into Θ which causes no further insertions in Σ since $\text{Any}(x) = \Lambda$ for all x . Attempting to insert $\langle \text{Integer} \rightarrow, \text{Any} \rightarrow \rangle$ into Σ causes us to insert $\langle \text{Any}, \text{Integer} \rangle$ into Θ , which in turn causes us to attempt to insert $\langle \text{Any}(+), \text{Integer}(+) \rangle$ into Σ which fails since $\text{Any}(+) = \Lambda$ and $\text{Integer}(+) \neq \Lambda$. Thus we conclude that IntegerStack does not conform to AnyStack .

Note that there is no need to start with empty relations Θ and Σ ; any valid relation Θ on types and its induced relation Σ on signatures may be used as a starting point. In actually implementing this procedure, the relations Θ and Σ may be retained after conformity checking, thus eliminating the need to recompute them.

E. Comparison to Other Languages

One currently popular view of a data type is as a set of operations. That is, each data type is a set of operations that may be applied to values of the type, thereby providing an interpretation for those values [11], [23]. There are, however, two approaches to the question of the ownership of these operations.

A number of languages, including Russell [9] and CLU [23], consider the operations to be *fields* of the type; that is, the operations are owned by the type, and are selected from the type. As an example, consider the type *Integer*. In Russell or CLU, the type would own operations like:

```
operation 0 → Integer
operation + (Integer × Integer) → Integer
operation * (Integer × Integer) → Integer
```

In order to add two integers in CLU (Russell is similar), one selects the $+$ operation from the *Integer* type to apply to the values:

```
result := Integer$(argument1, argument2)
```

The other approach is exemplified by Smalltalk. In Smalltalk, the operations that may manipulate objects are broken into two disjoint sets. Some operations, generally those that create new objects, are defined as messages to (operations on) a distinguished object called the *class*. An example of a class operation is creating an integer value from a character string representing the value. The other set of operations are defined as messages to the object itself; the operations $+$ and $*$ on integers are examples. The Smalltalk statement that adds two integers and assigns the result to a variable is:

```
result ← target + argument
```

This statement is interpreted to mean “send the message with the name $+$ to the object *target*, providing the object *argument* as argument, and assign the resulting object to the variable *result*.”

In Emerald, we have adopted an approach similar to Smalltalk’s. Our primary motivation was a concern over distribution. In the CLU and Russell approach, the statements in the body of an operation have access to the representation of an arbitrary number of objects of the type of interest. In particular, the implementor of the addition operation on integers is allowed to access the bits that make up the representation of *both* of the arguments. In the face of distribution, implementing this is potentially very complicated. Either the two objects need to be moved to a common location in order to perform the operation, or else the implementation needs to be able to perform arbitrary primitive operations remotely. In our view, only the object that is performing the operation has its representation accessible, and all other objects, including others of the same type, must be accessed through invocations.

The view of operations as being owned by objects is also required by our separation of abstract types and implementations. Since the two integers being added may in fact be implemented differently, the only way of implementing the integer operations is by invoking lower level operations that each concrete type may implement differently.

F. Type Implementation Considerations

A major consequence of the flexibility of Emerald’s type system is that multiple objects are allowed to implement a single abstract type. In principle, multiple implementations of a data type cause no problems because operations are invoked by name, and it is up to the invoked object to understand the name. In practice, it is important to optimize this process.

In Smalltalk and Eden, the cost of invocation by name appears at *invocation time* since the appropriate code for each invocation must be located when the invocation is attempted. Substantial effort has been put into implementations of Smalltalk that reduce the overhead of code location (*method lookup* in Smalltalk terminology) to an acceptable level [10], [21]. Techniques include caching of operation addresses at call sites and making likely guesses about the implementations of objects with certain abstract types (like *Integer*), thereby optimizing for the common case. Since our object model is similar to that of Smalltalk, all of the optimizations that can be done in the implementation of Smalltalk could also be done in Emerald. In addition, the fact that Emerald is statically typed allows an alternative code location strategy that may yield better performance.

With each identifier, we associate a vector of operations. The length of this vector is determined by the identifier’s abstract type, since only the operations defined by

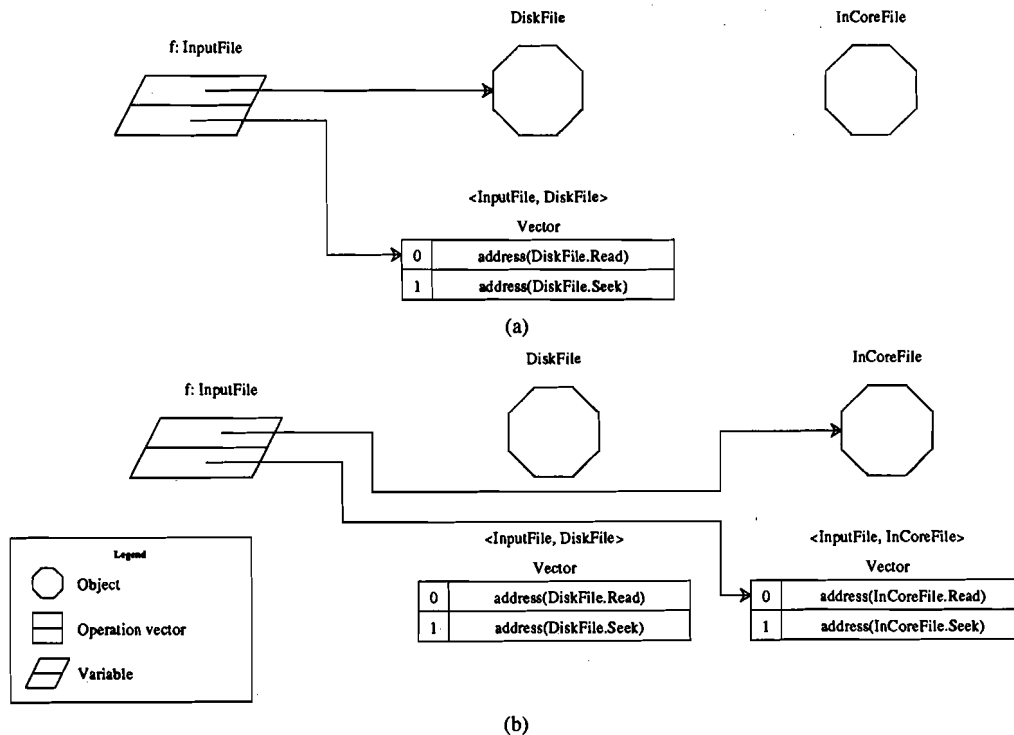


Fig. 3. (a) Before the assignment. (b) After the assignment.

the abstract type may be invoked on the object named by the identifier. The content of the vector is determined by the implementation of the object currently named by the identifier, since it consists of the addresses of the machine code sequences that implement the operations. Consider a variable f of abstract type a . The operations that may be performed on f are determined by a ; the addresses of the appropriate code for these operations depend on the implementation c of the object named by f . Thus the pair $\langle a, c \rangle$ uniquely determines an *operation vector* associated with f . In Fig. 3(a), a is *InputFile* and c is *DiskFile*. The vector has one element for each of the *InputFile* operations *read* and *seek*; the values of these elements are the addresses of the corresponding *DiskFile* routines.

When an assignment is made to f , the contents of its operation vector may need to be changed. For example, when an *InCoreFile* object is assigned to f , the operation vector associated with f must become appropriate to the pair $\langle \text{InputFile}, \text{InCoreFile} \rangle$, as shown in Fig. 3(b).

This scheme replaces the method lookup required by Smalltalk by a single indexing operation. The cost is an additional word of storage for the pointer to the operation vector, and occasional recomputation of the elements of these vectors on assignment. The operation vectors themselves may be shared between all identifiers of identical abstract type that name objects with the same implementation, since it is the pair $\langle \text{abstract type}, \text{implementation} \rangle$ that determines the contents of the vector.

V. DISTRIBUTION SUPPORT

As previously stated, the principal objective of Emerald is to simplify the construction of distributed programs. System concepts such as concurrency, multiple nodes, and

object location are integrated into the language. This differs from, for example, EPL, where distribution is layered on an existing language through the use of a preprocessor, and from Accent [27], where distribution is provided as an operating system facility.

In Emerald, objects encapsulate the notions of process and data and are the natural unit of distribution. At any time each Emerald object is located at a specific node. Conceptually, a node is an object of a system-defined type. Node objects support node-specific operations, thereby allowing objects to invoke kernel operations. Such access to the underlying kernel is analogous to that provided by *kernel ports* in Accent.

Programmers may choose to ignore or exploit the concept of object location. In a distributed system, objects must be able to invoke other objects in a location-independent manner. This facility makes network services transparently accessible. In Emerald, locating the target of an invocation is the responsibility of the system. An object is permitted to move between successive invocations, or even during an invocation. While applications can control the placement of objects, most applications can ignore location considerations since the semantics of local and remote invocation are identical.

Nevertheless, there are two reasons for making location visible to the programmer: performance and availability. In a network, the efficiency of interobject communication is obviously a function of location. An application can colocate objects that communicate intensely and thus reduce the communication overhead. Alternatively, numerical applications can achieve significant performance gains by placing concurrent subcomputations on different nodes. An object manager may increase availability by placing replicas of its objects on different nodes.

A. Object Mobility

An important feature of the Emerald design is support for unrestricted object mobility. Mobility (sometimes called *process migration* in nonobject systems) has been discussed in the literature as a design goal for many operating systems [1], [7], [26], [27], [34], although full implementations are not common. Mobility can be used to improve the overall performance of a distributed system by load sharing, i.e., moving objects from loaded nodes to idle nodes. For example, when a user starts a compilation and concurrently continues editing, it is often an advantage to move the compilation to another (less loaded or more powerful) node [12]. The advantages of load sharing are offset by the cost of obtaining load information and the cost of moving objects. The lower the cost of object relocation, the greater the potential for performance improvements. Therefore a goal of Emerald is to provide low-cost mobility.

Emerald differs from previous systems with migration facilities in the grain of mobility. Systems such as V [34], DEMOS/MP [7], and LOCUS [26] use processes (and their associated address spaces) as the unit of mobility. Emerald advocates the use of a single object model and therefore allows any object, even a single data item, to be the unit of migration.

Emerald provides two primitives to control location. A programmer can *fix* the location of an object at a specific location, or the programmer can *move* an object to a specific location. The location is specified either directly, by referring to a node object, or indirectly, by referring to another object's current location. Thus, while programmers can ignore location with respect to invocation, they can also choose to benefit from the advantages of specific locations within the network.

B. Parameter Passing and Call-by-Move

In Emerald, objects communicate through invocation. Since the invocation parameters must themselves be objects, the natural parameter passing method is call-by-object-reference. This is similar to CLU and Smalltalk. However, call-by-reference does present a problem when used in a distributed system. On a remote invocation, references by the called object to its parameters are likely to require remote invocations also. To avoid seriously degrading performance, systems such as Argus have assumed that parameters to remote calls must be passed by value [17]. In Emerald, because objects are mobile, it is possible to avoid most remote references to invocation parameters by moving them to the site of the callee. Whether or not this is worthwhile depends on the size of the parameter object, the number of active invocations of the object, and the number of invocations of the object to be initiated by the destination.

Argument motion may be appropriate in two cases. First, based on type information available at compile time, the compiler may choose to move an object along with the invocation. For example, since integers are small and immutable, they are obvious candidates for relocation.

Second, the programmer may explicitly request that an object be moved with the invocation, based on his own information. This is done by means of a new parameter passing mode that we call *call-by-move*. A call-by-move parameter object is passed by reference (as is any other parameter). However, at call time the call-by-move argument object is moved to the destination node, allowing it to be efficiently referenced by the invokee. The programmer can also specify that the argument is to move back upon completion of the invocation.

Call-by-move is a performance optimization, since it could be emulated by first moving each of the call-by-move parameter objects to the invokee's node, and then invoking the object. However, performing the moves separately would cause multiple messages to be sent across the network. For small objects, it is significantly more efficient to package the referenced objects together with the invocation, thereby reducing the network traffic and message count.

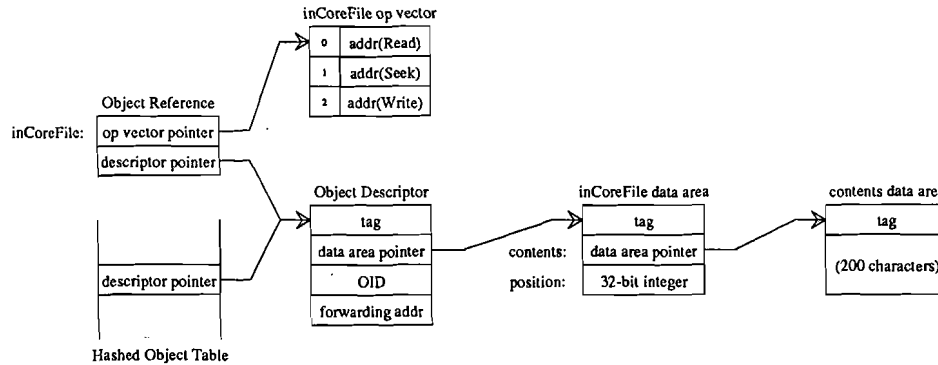
Although call-by-move reduces the cost of references made by the invokee, it increases the cost of the invocation itself. If the parameter object is mutable and shared, it also increases the cost of references by the invoker. One of the goals of Emerald is to enable us to investigate these tradeoffs experimentally.

C. Addressing and Storage Structure

Performance is an important goal of the Emerald system, and our implementation stresses a high degree of integration between the compiler and the run-time kernel. The interface between them is defined so that many functions can be performed either by the kernel, by the compiler, or by compiler-generated code. For example, the compiler may detect that an invocation will always be local and generate code that circumvents the more general invocation mechanism of the kernel. Thus, while Emerald has a uniform object model, objects can be implemented in different ways using different addressing mechanisms, different storage structures, and different invocation protocols.

As previously stated, simple local objects can be directly allocated within the containing object's data area. Such objects are addressed directly by offset in the data area; no other structures are needed. More complex local objects are heap allocated and are accessed through a 32-bit data address stored in the object's data area. Finally, remotely accessible objects have a 64-bit *object reference* in the containing object's data area. An object reference consists of the address of an *object descriptor* and the address of the object's operation vector. The object descriptor contains necessary control information including the object's current location and a pointer to the object's data area.

To support object mobility, internode object references either must be location independent or must be translated when crossing node boundaries. Both options impact performance; however, the impact is lessened by integrating distribution support into the programming language. We have chosen to optimize for the local case at a small pen-

Fig. 4. Storage layout for `inCoreFile`.

ality to mobility. Therefore, efficient location-dependent references are used locally, but these references must be translated when an object moves.

In addition to an object descriptor, each remotely accessible object is assigned a unique *object id* (OID). Each node has an *Object Table* that stores pointers to the local object descriptors for all remotely accessible objects. The object table is accessed by a hash function using the OID as a key. When a reference to the object is sent across the network, the OID and a *forwarding address* (see Section V-E) for the object are sent along as well. At the destination node, the reference is translated to a node-local object descriptor address by looking up the OID in the object table. If the lookup fails then a local object descriptor is created using the OID and forwarding address passed with the reference.

If an object moves then only the descriptor need be changed. This allows for efficient referencing of node-local objects—there is no need for a table lookup as is the case for referencing Accent ports [27], Eden objects [1], or Smalltalk objects [15].

For example, Fig. 4 shows the addressing and storage structure for an `inCoreFile` object as previously defined in Fig. 1. On the left of the figure are a reference to the `inCoreFile` object and the Object Table. The `inCoreFile` reference contains pointers to the object descriptor and an operation vector. The `inCoreFile`'s object descriptor points to its data area. Inside of the data area are two variables: *contents*, a reference to a vector object, and *position*, a reference to an integer object (lines 6 and 7). The vector object has been heap allocated and is referenced by its heap address. The integer object has been allocated directly in the `inCoreFile`'s data area. In this case, because both objects are local and both implementations are known at compile time, the compiler does not generate an operation vector or procedure calls; instead, it generates in-line machine instructions to execute these objects' operations. For example, the addition and assignment in line 11 of Fig. 1 reduces to a single VAX increment instruction, while the `getElement` invocation on line 10 produces in-line instructions to perform the bounds check and move. Objects that may be referenced remotely are also heap allocated but would be referenced indirectly through an object descriptor, as noted above.

D. Moving Objects

To support object mobility, the Emerald compiler generates relocatable code and provides the kernel with descriptions of the structure of objects. These compiler-generated descriptions, called *templates*, enable the kernel to find the variables that need to be translated upon migration. To move an object, the object data area is traversed using the compiler-generated template. During the traversal, a *translation table* is constructed that maps node-local addresses of global objects into OID's. Local objects referenced by the migrating object are moved together with the object data area and their node-local addresses are also entered into the translation table. At the destination, the object data area is rebuilt and the references are translated into node-local addresses.

The migrating object may have one or more processes executing within it. These processes are suspended and the activation records for operations of the object are moved along with the object data area. In this way, the stack becomes segmented; each segment is represented internally as if it were an object, i.e., it is given an object descriptor and an OID. When an operation returns, the stack segment to return to may be found using the general object location algorithm (described in the next section).

The precopy employed by the V system [34] avoids suspending processes during most of the migration. Since most objects in our system will be significantly smaller than the logical hosts of V, we would not expect to gain much from using this precopy technique. Furthermore, the technique requires access to the dirty bits of the paging hardware and such access is not available to our current prototype.

E. Keeping Track of Mobile Objects

When an object moves, it leaves behind a *forwarding address* in the object descriptor so that invocations sent to its old location can be forwarded. Every time a reference is passed between nodes, the forwarding address is passed as well. To enable a node to decide which of two forwarding addresses is the most recent, a simple *timestamp* is included in the forwarding address. Fowler [13], [14] has shown that it is sufficient to let the timestamp merely be the number of times an object has attempted to move. We plan to experiment with several of the for-

warding address propagation policies outlined by Fowler and with the policy used for updating DEMOS/MP [26] links.

When machines crash, forwarding addresses are lost unless they are kept on stable storage. Emerald is designed to operate within a local area network; this makes it possible to recover forwarding addresses by exhaustive search, and thus avoids committing forwarding addresses to stable storage on every update. The exhaustive search is implemented using a *reliable broadcast* protocol [8] which is optimized for the case when the referenced object is accessible.

VI. CURRENT STATUS AND PERFORMANCE

Our current prototype of Emerald and its run-time system runs on a small local area network of DEC MicroVAX II computers running ULTRIXTM (DEC's version of Berkeley UNIXTM 4.2 bsd). Using UNIX avoids the need for implementation of low-level machine-specific and device-specific software. The cost of UNIX is seen in the performance of certain functions, most notably machine-to-machine communications. Our use of UNIX does not have significant impact on the implementation of objects within a single node.

Each node is implemented as a UNIX process and all objects residing on a particular node share that single address space. As in Concurrent Pascal [5] and Mesa/Pilot [32], protection is provided by the compiler; no hardware protection scheme is used. The prototype does not support multiple languages—only Emerald programs may be run on top of the Emerald kernel. Processes are lightweight, making context switching fast; processes and monitors are implemented in a manner similar to Concurrent Pascal [6]. The use of a single address space makes it possible to implement process management and node-local communication efficiently. Since objects can move to and from a node, dynamic loading is supported. This also enables the kernel to migrate an object to backing store so that its data area can be reused.

The Emerald compiler and run-time system are currently able to support the creation and invocation of Emerald objects on a single processor or across the network. We are currently working on the implementation of object mobility and distributed garbage collection. The performance of local operations is summarized in Table II. The implementation of remote invocations is not yet at a stage where performance figures are meaningful.

VII. SUMMARY

We have described the approach to typing and distribution taken in Emerald. By supporting abstract types, Emerald recognizes that a given abstraction will often be implemented in several ways. This allows a system to evolve and permits the coexistence of multiple compiler-generated implementations that take advantage of distri-

TABLE II
PERFORMANCE OF LOCAL OPERATIONS

Example	Implementation	time (μ s)
$position \leftarrow position + 1$	inline call	1.5
$c \leftarrow contents.getElement(position)$	inline call	6.3
$inCoreFile.Seek[0]$	potentially remote call	20.8

bution and locality. In addition, the concept of object mobility and call-by-move allows call-by-reference semantics to be used uniformly at all interobject interfaces, thus simplifying the construction and reconfiguration of distributed applications.

We believe that the integration of these *system* concepts into a programming language will allow us to use the same object model for both programming in the small and for programming in the large, and to achieve an efficient implementation in both cases. We are building a prototype compiler and run-time system to demonstrate the viability of our approach.

REFERENCES

- [1] G. T. Almes, A. P. Black, E. D. Lazowska, and J. D. Noe, "The Eden system: A technical review," *IEEE Trans. Software Eng.*, vol. SE-11, pp. 43-59, Jan. 1985.
- [2] A. D. Birrell and B. J. Nelson, "Implementing remote procedure calls," *ACM Trans. Comput. Syst.*, vol. 2, no. 1, pp. 39-59, Feb. 1984; presented at the Ninth ACM Symp. Operat. Syst. Principles, Oct. 1983.
- [3] A. P. Black, "The Eden programming language," Dep. Comput. Sci., Univ. Washington, Seattle, Tech. Rep. 85-09-01, Dec. 1985.
- [4] A. Black, N. Hutchinson, E. Jul, and H. Levy, "Object structure in the Emerald system," in *Proc. ACM Conf. Object-Oriented Program. Syst., Lang. and Applications*, Portland, OR, Oct. 1986, pp. 78-86.
- [5] P. Brinch Hansen, "The programming language Concurrent Pascal," *IEEE Trans. Software Eng.*, vol. SE-2, pp. 199-205, June 1975.
- [6] —, *The Architecture of Concurrent Programs*. Englewood Cliffs, NJ: Prentice-Hall, 1977.
- [7] D. A. Butterfield and G. J. Popek, "Network tasking in the LOCUS Distributed UNIX system," in *Proc. USENIX Summer 1984 Conf.*, 1984.
- [8] J. Chang and N. F. Maxemchuk, "Reliable broadcast protocols," *ACM Trans. Comput. Syst.*, vol. 2, no. 3, pp. 251-273, Aug. 1984.
- [9] A. Demers and J. Donahue, "Revised report on Russell," Dep. Comput. Sci., Cornell Univ., Ithaca, NY, Tech. Rep. 79-389, Sept. 1979.
- [10] L. P. Deutsch and A. M. Schiffman, "Efficient implementation of the Smalltalk-80 system," in *Proc. Eleventh ACM Symp. Principles of Program. Lang.*, Salt Lake City, UT, Jan. 1984, pp. 297-302.
- [11] J. Donahue and A. Demers, "Data types are values," *ACM Trans. Program. Lang. Syst.*, vol. 7, no. 3, pp. 426-445, July 1985.
- [12] D. L. Eager, E. D. Lazowska, and J. Zahorjan, "Dynamic load sharing in homogenous distributed systems," *IEEE Trans. Software Eng.*, vol. SE-12, pp. 662-675, May 1986.
- [13] R. J. Fowler, "Decentralized object finding using forwarding addresses," Ph.D. dissertation, Dep. Comput. Sci., Univ. Washington, Seattle, Tech. Rep. 85-12-1, Dec. 1985.
- [14] —, "The complexity of using forwarding addresses for decentralized object finding," in *Proc. Fifth SIGACT/SIGOPS Conf. Principles of Distributed Computation*, Calgary, Alta., Canada, Aug. 1986.
- [15] A. Goldberg and D. Robson, *Smalltalk-80: The Language and its Implementation*. Reading, MA: Addison-Wesley, 1983.
- [16] I. Greif, R. Seliger, and W. Weihl, "Atomic data abstractions in a distributed collaborative editing system," in *Proc. Thirteenth ACM Symp. Principles of Program. Lang.*, Jan. 1986.
- [17] M. Herlihy and B. Liskov, "A value transmission method for abstract data types," *ACM Trans. Program. Lang. Syst.*, vol. 4, pp. 527-551, Oct. 1982.
- [18] C. A. R. Hoare, "Monitors: An operating system structuring concept," *Commun. ACM*, vol. 17, no. 10, pp. 549-557, Oct. 1974.

TMULTRIX is a trademark of Digital Equipment Corporation.

TMUNIX is a trademark of AT&T Bell Laboratories.

- [19] R. C. Holt, *Concurrent Euclid, the Unix System, and Tunis*. Reading, MA: Addison-Wesley, 1983.
- [20] A. K. Jones, R. J. Chansler, Jr., I. Durham, K. Schwans, and S. R. Vegdahl, "StarOS, a multiprocessor operating system for the support of task forces," in *Proc. Seventh ACM Symp. Operat. Syst. Principles*, Dec. 1979, pp. 117-127.
- [21] G. Krasner, Ed., *Smalltalk-80: Bits of History, Words of Advice*. Reading, MA: Addison-Wesley, 1983.
- [22] B. Liskov, "Primitives for distributed computing," in *Proc. Seventh ACM Symp. Operat. Syst. Principles*, Dec. 1979, pp. 33-42.
- [23] B. Liskov, R. Atkinson, T. Bloom, E. Moss, C. Schaffert, B. Scheffler, and A. Snyder, "CLU reference manual," Lab. Comput. Sci., Massachusetts Inst. Technol., MIT/LCS/Tech. Rep.-225, Oct. 1979.
- [24] B. Liskov and R. Scheffler, "Guardians and actions: Linguistic support for robust, distributed programs," in *Proc. Ninth ACM Symp. Principles of Program. Lang.*, 1982.
- [25] B. Liskov, "Overview of the Argus language and system," Lab. Comput. Sci., Massachusetts Inst. Technol., Programming Methodology Group Memo 40, Feb. 1984.
- [26] M. L. Powell and B. P. Miller, "Process migration in DEMOS/MP," in *Proc. Ninth ACM Symp. Operat. Syst. Principles*, Oct. 1983, pp. 110-119.
- [27] R. F. Rashid and G. G. Robertson, "Accent: A communication oriented network operating systems kernel," in *Proc. Eighth ACM Symp. Operat. Syst. Principles*, Oct. 1981, pp. 64-75.
- [28] J. C. Reynolds, "Using category theory to design implicit conversions and generic operators," in *Semantics-Directed Compiler Generation (Lecture Notes in Computer Science)*. New York: Springer-Verlag, 1980.
- [29] J. C. Reynolds, "The essence of Algol," in *Algorithmic Languages*, J. W. de Bakker and J. C. van Vliet, Eds. Amsterdam, The Netherlands: North-Holland, 1981.
- [30] C. Schaffert, T. Cooper, B. Bullis, M. Kilian, and C. Wilpolt, "An introduction to Trellis/Owl," in *Proc. ACM Conf. Object-Oriented Program. Syst., Lang., and Applications*, Portland, OR, 1986, pp. 9-16.
- [31] R. W. Scheffler, "An analysis of inline substitution for a structured programming language," *Commun. ACM*, vol. 20, no. 9, pp. 647-654, Sept. 1977.
- [32] R. E. Sweet, "The Mesa programming environment," in *Proc. ACM SIGPLAN 1985 Symp. Program. Lang. Issues in Program. Environments*; also *SIGPLAN Notices*, vol. 20, no. 7, pp. 216-229, June 1985.
- [33] W. Teitelman, "A tour through Cedar," in *Proc. Seventh Int. Conf. Software Eng.*, Orlando, FL, Mar. 1984.
- [34] M. M. Theimer, K. A. Lantz, and D. R. Cheriton, "Preemptable remote execution facilities for the V-system," in *Proc. Tenth ACM Symp. Operat. Syst. Principles*, Orcas Island, WA, Dec. 1985, pp. 2-12.
- [35] W. A. Wulf, R. Levin, and S. P. Harbison, *HYDRA/C.mmp: An Experimental Computer System*. New York: McGraw-Hill, 1981.



Andrew Black was born in London, England. He received the B.Sc. (Hons.) degree from the University of East Anglia in computing studies, and the D.Phil. degree from the Programming Research Group of the University of Oxford in programming languages and software engineering.

He has been on the faculty of the Department of Computer Science at the University of Washington, Seattle, since 1981. His current research interests are in the areas of distributed systems and programming language design, and in the way

these areas interrelate.

Dr. Black is a member of the Association for Computing Machinery, the British Computer Society, the Union of Concerned Scientists, and Computer Professionals for Social Responsibility, and an affiliate of the IEEE Computer Society.



Norman Hutchinson received the B.Sc. degree in computer science from the University of Calgary, Calgary, Alta., Canada, in 1982, the M.S. degree in computer science from the University of Washington, Seattle, in 1985, and expects to receive the Ph.D. degree from the University of Washington in late 1986. His thesis topic is programming language design for distributed systems.


Effective January 1987, he will be an Assistant Professor of Computer Science at the University of Arizona in Tucson.



Eric Jul was born in Roskilde, Denmark. He received the Cand. Scient. (M.S.) degree from the University of Copenhagen, Copenhagen, Denmark, in 1980.

Currently, he is working toward the Ph.D. degree at the University of Washington, Seattle. His research interests are object-oriented systems, the design and implementation of programming languages for distributed systems, and distributed garbage collection. His thesis topic is object mobility in distributed systems.

Mr. Jul is a member of the Association for Computing Machinery, Computer Professionals for Social Responsibility, and the Danish Computing Society.



Henry Levy received the B.S. degree from Carnegie-Mellon University, Pittsburgh, PA, and the M.S. degree from University of Washington, Seattle.

He is Research Assistant Professor in the Department of Computer Science at the University of Washington and a Consulting Engineer on leave from Digital Equipment Corporation. His research interests include operating systems, distributed systems, and computer architecture. He is author of the book *Capability-Based Computer*

Systems and coauthor of *Computer Programming and Architecture: The VAX-11*.



Larry Carter received the A.B. degree in mathematics from Dartmouth College, Hanover, NH, in 1969, and the Ph.D. degree in mathematics from the University of California at Berkeley in 1974.

Since 1975, he has been at IBM's Thomas J. Watson Research Center in Yorktown Heights, NY, except for sabbatical leaves to teach at the Penn State University and at Berkeley, and extended visits to the University of Washington in Seattle and to Hampshire College in Amherst,

MA. His research interests include the theory of computation and VLSI.