

# A Dossier Driven Persistent Objects Facility

Robert Mecklenburg  
Charles Clark  
Gary Lindstrom  
Benny Yih

UUCS-94-002

Department of Computer Science  
University of Utah  
Salt Lake City, UT 84112 USA

January 10, 1994

## Abstract

We describe the design and implementation of a persistent object storage facility based on a dossier driven approach. Objects are characterized by dossiers which describe both their language defined and “extra-linguistic” properties. These dossiers are generated by a C++ preprocessor in concert with an augmented, but completely C++ compatible, class description language. The design places very few burdens on the application programmer and can be used without altering the data member layout of application objects or inheriting from special classes. The storage format is kept simple to allow the use of a variety of data storage backends. Finally, by providing a generic object to byte stream conversion the persistent object facility can also be used in conjunction with an interprocess communication facility to provide object-level communication between processes.<sup>1</sup>

---

<sup>1</sup> This research was sponsored by Hewlett-Packard’s Research Grants Program and by the Advanced Research Projects Agency (DOD), monitored by the Department of the Navy, Office of the Chief of Naval Research, under Grant number N00014-91-J-4046. The opinions and conclusions contained in this document are those of the authors and should not be interpreted as representing official views or policies, either expressed or implied, of the Defense Advanced Research Projects Agency, the U.S. Government, or Hewlett-Packard.



# 1 Motivation

The basic problem of a persistent object store (POS) is simply stated:

Given a reference to the root node of a graph of objects generate a data stream which can be used to reconstitute the original object graph at a later time.

Many approaches have been pursued to solve this basic problem (see Section 11 for a summary). The utility of these approaches is governed by the constraints they impose on application code in such dimensions as (i) language or compiler extensions, (ii) mandatory inheritance from library base classes, (iii) system transformation of application source code, (iv) expansion of object size, (v) mandatory presence of virtual function tables, and (vi) programmer declaration of supporting functions and observance of programming style restrictions.

We describe a new approach which poses no constraints in (i) – (v), and minor client obligations in (vi). Our approach is based on preprocessor-generated dossier objects[13], which drive fully polymorphic (i.e., applicable to all types) load and store functions. In addition to supporting object persistence, our approach provides a fully general means for transporting object graphs in address space independent form (i.e., “pickled”, with “unswizzled” pointers). Our design has been motivated by the stringent demands of a large (750,000 line) C++ CAD/CAM/visualization application[2].

## 2 What Is An Object?

We begin by defining our unit of persistence, which we term an *object*. While some approaches take this to be C++ class instances, this basis is too narrow for applications such as our CAD client, which make extensive use of graphs of vectors and structures, with semantically significant sharing relationships. Hence we define an *object* to be a contiguous region of memory whose type is known either through static type information, through dynamic type information (e.g., virtual function table), or information provided by the application programmer. An object is identified in an application by a pointer or reference to its first address along with some notion of its bounds (derived from type information). We explicitly disallow pointers to the interior of objects. An *object graph* consists of a collection of objects formed into an arbitrary graph by pointers embedded in the objects. An object is identified in the persistent store by a unique *object identifier* (OID). An application requests objects by OID and can access the OID of an object given its virtual address in the application.

## 3 Client Constraints

To be as convenient as possible a POS must minimize the impact of its use on application source code and the software development process while at the same time maximizing functionality. Among the features of a POS, we feel the following to be important: minimal impact on object layout and class declarations, allow the use of standard language tools, provide object access from a variety of hardware platforms, provide object access after class mutation. We discuss each of these requirements in turn.

The POS should not require “large” changes to class definitions. In particular, any system which requires altering the class layout by adding data members, virtual functions (where none existed before) or additional base classes is unacceptable. Such a system would impose storage overhead and incompatibilities which many applications could not tolerate. However, adding additional virtual functions to a class with an existing virtual function table would allow more convenient use of the storage facility. If this modification were allowed (but optional) it would provide for a convenient interface for application specific classes while still allowing library classes (for which there is no source code) to persist.

One of the biggest stumbling blocks for POSs is the requirement for non-standard language tools (e.g., special compilers) to enable objects to persist. These tools either parse an extended language syntax (translating into standard C++) or generate augmented class implementations (or both). Our group, having worked on large software projects using these approaches, find them burdensome; chose to require the class definition be written in standard C++. This means that there is only one class definition (with no additional semantic information in other files) and that applications can be compiled and run (albeit without persistence) with or without the persistent objects facility. This significantly simplifies porting and piece-wise development and testing of applications.

Once a POS is integrated into an application or organization its use quickly becomes fundamental to the project and the persistent objects themselves become a valuable resource. As such, it is often unacceptable to abandon the database when new hardware or software is acquired or when class definitions change. Furthermore, as the size of the database grows evolving the data *en masse* becomes a significant burden. We feel a more reasonable approach is to integrate platform heterogeneity and type evolution cleanly into the persistent store allowing for lazy transformation of objects to the reader's requirements.

We discuss other, less major, constraints on the POS as they arise.

## 4 An Object Description Language

Next, we address the need for a language in which to describe objects. An object which is an instance of a primitive C++ type may be described simply by its standard type name. One may reasonably expect that an object which is an instance of a class may be described by the C++ declaration of that class. Indeed, to a first approximation, that is correct. Unfortunately, there are several "extra-linguistic" patterns of use which are not sufficiently described by standard C++ syntax, particularly with respect to dynamically sized objects (e.g., strings and other vectors). The problem is to identify important idioms required by applications and to provide an annotation mechanism which does not invalidate the use of standard language tools. In addition to these annotations, the POS may require classes to provide various semantic handles to allow storage and retrieval.

The most important idiom in C++ which is not adequately described by class declarations is the use of pointers to access dynamically sized regions of memory. Strictly interpreted, the declaration:

```
char *path;
```

identifies a pointer to an unknown number of characters. By convention the number of characters is determined by a *sentinel value*, in this case the null character. The sentinel value technique for dynamically sized data can be used with any data type, but is most typically used with pointers and integral types where the zero bit pattern is used as the sentinel. A competing style for identifying the size of dynamically sized memory regions relies on a pair of data values:

```
int    n;    // size of name
char *name;
```

where the dynamic size is stored explicitly in a separate data member.

Static data members of a class pose a different sort of problem for a POS. Indeed, one may question whether static data members should persist at all. Often these data members are used to resolve issues inherent in run-time data management. For instance, an application might maintain an *extent* list of all allocated instances. Such a list acquires a completely different meaning in a persistent store owing to the shared, distributed, and concurrent nature of the store. Our approach is to store static data members, but not to manage concurrent access. Aside from ensuring consistent concurrent writes for single data members we do not assume any further capabilities of the underlying POS such as notifying readers of

updates to shared data. Similar to static data members there may be non-static data members which the programmer does not want saved. For example, an object might contain a pointer to a buffered file structure which has no meaning (or a different meaning) when stored in a POS. These nodes can be annotated as *orphaned* objects; their value will not be stored and their pointers will not be traversed.

How can these annotations be applied to a class definition if standard compilers are used and no additional files are consulted? There are two basic approaches possible: embedded annotations in comments and augmented identifier names. The first approach places comments adjacent to data members containing keywords identifying various attributes. The second approach uses the data member name itself (or its type name) to contain the attribute. An example of this might be:

```
typedef char char__null; // Null terminated string.
char__null *path;

typedef int int__sized; // Integer sized string.
int__sized n;
char *      name;
```

We chose this technique for several reasons: it allows the dossier generator to use the C preprocessor (which elides comments), it does not interfere with a standard commenting style for class declarations, and it allows us to experiment with a novel annotation technique. Furthermore, annotating the type of the data member (rather than the member itself) leaves the application programmer free to select meaningful member names unencumbered by the annotations. The currently supported annotations are:

```
__null    dynamically sized, zero terminated
__sized   dynamically sized, this member is the size, following member is the pointer
__orph    an orphaned object, don't save
```

The final class of semantic problems we discuss relates to handling application or environment specific meanings associated with objects. Examples of such problems include storing hash tables and file handles. As with other members the writer of the object must annotate the stored instance with information allowing the reader to reconstitute a similar object with semantics equivalent to the original object. For a hash table, the reader may have a different hash function or table size and therefore must rehash the members of the table. For a file handle, the reader must find and open the file and set the current position. An annotation on a declaration cannot transmit this information (and indeed, may not have the information to transmit). To allow for this type of application specific behavior the application programmer can define load and store *hooks* which are called by the POS during object I/O. The load/store hook has a special name and type signature recognized by the dossier generator:

```
void __load_store_hook( int when );
```

This member function is added to the class declaration of any class requiring special handling during I/O. The function can be called under three circumstances (indicated by the **when** parameter): after loading an object, before storing an object, and after storing an object.

When an object is restored from the POS several application and implementation specific initializations must be performed. The most obvious of these is setting the virtual function table pointer. This can be done in a variety of ways: from using the **new** placement syntax and having the application programmer invoke the constructor to copying the pointer from an initialized sample instance. The later approach does not allow for the application to gain control during object allocation and is therefore unacceptable. Using the **new** placement syntax has the problem of compatibility with other software packages (including the application's classes). A compromise requires the application class to define a special constructor which we call the *reconstructor*. This approach allows classes to overload

`new` and `delete` and to gain control during object construction. The reconstructor is identified by its type signature:

```
<class_name>( reconstructor_t );
```

Finally, to allow convenient use of the POS with polymorphic objects we encourage the application programmer to declare a virtual function for accessing the dossier of a class:

```
virtual dossier_c *__get_dossier() const;
```

This allows the application and POS interface to access the dossier of conforming objects simply. For objects which do not support the `__get_dossier` member function, the application must provide the dossier handle explicitly. This results in a simple and convenient interface for classes under application programmer control, while still allowing other classes to persist. Once the dossier for the root object is obtained, dossiers for other objects in the graph can be accessed through the root object dossier.

Once an application's class declarations (e.g., `.h` files) have been adapted to express these extra-linguistic features, they become the application's class description. These files are read and analyzed by a preprocessor based on the C++ grammar (written by James Roskind[22]). The preprocessor emits auxiliary C++ files which construct instances of class dossiers embodying the class descriptions, including associated annotations. These emitted files are compiled and linked, along with a support library, into an application to implement the client side of the POS. Note that client source files are only read, not transformed, in this process.

## 5 Capture of Compiler and Platform Characteristics

To build a complete description of objects, including data member layout, the dossier generator must mirror the algorithms of the current compiler and would therefore not be particularly portable. We avoid this problem by separating the dossier into machine/compiler independent and dependent portions. The compiler independent portion is constructed by the dossier generator while the dependent portion is computed at run-time from auto-configuring code written into the dossier initializer. The compiler and machine dependent structures gather three types of information: size and format of data types, location of data members in objects, and handles on member functions. We discuss each briefly.

To allow dossier code to read and write objects on differing platforms (both hardware and software) the polymorphic I/O code must know the size of each data type and its format when written to a persistent store. Size information is easily acquired through the use of the `sizeof` compiler directive. Also, byte order and floating point format must be determined. In the worst case, these characteristics must be explicitly specified for each platform making the dossier source code non-portable. In the normal case, however, byte order can be determined through simple calculations and IEEE standard floating point format can be assumed.

The location of data members and base classes for an object are determined using a technique similar to the ANSI C `offsetof` macro. For each (non-static) data member, its location is determined by taking its address and subtracting the object's base address. This requires that the dossier initializer be either a friend or member function of the class. Base class offsets are calculated similarly by casting a "pointer to derived class" to a "pointer to base class". For example, if class `D` derives from class `B`, the expression:

```
((B *)((D *)8)) - 8
```

returns the offset of a `B` within a `D`. (The use of a non-zero base address subverts optimizations in various compilers.) This expression is portable across all platforms (that we are aware of)[9].

Finally, the polymorphic I/O operations must invoke class reconstructors and load/store hooks to perform their functions. Since the address of a constructor cannot be computed, we wrap the reconstructor in a simple C++ function and store its address in the dossier. For uniformity we use the same technique to store the load/store hook in the dossier.

## 6 The Storage Algorithm

The basic storage algorithm is a simple graph traversal driven by the graph's root object and the dossiers. We begin by retrieving the OID of the object to be saved. If it does not have one, we allocate an OID. Then we place the object and its OID into the queue of objects waiting to be processed. The rest of the algorithm proceeds as follows:

### Algorithm 1

```
dequeue the next node to process
if the node is unsaved
    run the pre-store hook
    mark the object as saved
    enqueue all embedded pointers (allocate OIDs, if necessary)
    store the dossier, if necessary
    store the object and dossier OIDs, and machine id
    store the object
    store the OID of the target of every embedded pointer
    run the post-store hook
```

Dossiers are just objects so they are stored, along with the objects they describe, using the same algorithm. Of course, only one copy of the same dossier is stored and that dossier is referenced by all instances of that class through the OID of the dossier. Since a dossier is an object it must have a descriptor, or *meta-dossier*, to be read and written. This meta-dossier is a permanent component in the support library and is never written to or read from a POS or communication channel. The meta-dossier is generated by running the dossier generator over its own data structures.

The storage format is designed to be “retargetable” to different object storage engines and is therefore a mix of low-level formats and high-level information. The storage engines currently in use are a transactional DBM and a simple Unix file interface (an Exodus interface is planned). Writing is performed in the simplest possible way, by copying the machine representation of each data member value to the POS. It is the responsibility of the reader to decipher the writer's format. Since objects are often read and written on a single platform this proves reasonably efficient for local communication and temporary storage.

Retrieving object graphs is similar. The retrieval is initiated by the application with the OID of the root node of an object graph. This node is entered into a queue of nodes yet to be read and proceeds as follows:

### Algorithm 2

```
dequeue the next node to process
if the node is not yet read
    load the dossier of the object
    load the binary image of the object
    invoke the reconstructor to allocate memory for the object
    record the new object's address and OID
    copy the values of data members from the binary image to the new object
    for each pointer member set the new address, if available
```

```

        if not available, place pointer member on patch queue
    run the post-load hook
else
    return the address of the object
traverse patch queue, setting remaining pointer members

```

The object is loaded as a set of binary values from the original object. The dossier is used to pick through this bag of bits to identify data members and their values. The new values for pointers are accessed by the OID of the target object. Due to cyclic graph structures some objects will not have been read yet, so pointers to these objects must be queued until the desired object has been read.

## 7 Heterogeneity

Heterogeneity is handled by providing a machine description object which contains information concerning hardware and compiler specific data. In Algorithm 1 a machine identifier is stored along with the OIDs of the object and its dossier. This machine identifier references a structure describing the hardware characteristics (e.g., byte order, floating point format) and software characteristics (e.g., member layout) of the writer. When the data for an object is copied from the binary image of the writer to the run-time memory allocated for the reader machine dependent translations are performed.

Although the translations from one hardware platform to another must be hand-crafted, the actual process of converting values from one format to the other is controlled through the dossiers. To avoid writing  $n^2$  conversion routines a standard intermediate format can be used to reduce the number of conversion routines to  $2n$ .

## 8 Object Evolution

Invariably, the classes for objects stored in the POS will change due to changes in the user's requirements and added functionality. It is important that old data continue to be accessible to current applications. There are three basic approaches to evolving an object instance from one class declaration to another:

1. provide accessor functions,
2. copy using a "static" algorithm,
3. copy using a "dynamic" algorithm.

The first technique requires that an application be enhanced with accessors that know the old and new type and offset of the desired data member. This accessor is invoked on the old object and returns a value as if from a new object. This is unsuitable for many applications due to its highly hand-crafted nature. The second technique uses the dossier of the old and new objects to copy data member values one by one from the old to the new object using some fixed algorithm. Types that have changed may be converted if the conversion is sufficiently simple (e.g., int to float) and discarded otherwise (assuming that the old value has no translation). New data members may be initialized to some default value (e.g., zero). Experience with one large project indicates that this is a useful evolution technique for many simple object transformations[15]. Nevertheless, it is insufficient as the only (or even primary) type evolution mechanism. The final technique allows the application programmer to provide a function to translate an object from one version of a class to another.

Dossiers can be annotated with version information and translation functions capable of converting from one version of an object to another. The dossier driven type evolution system can then chain



conversion functions to evolve from one version of an object to the next until the desired version has been computed. A mixture of the second and third techniques described above is being implemented for our POS.

## 9 Current status

The dossier generator is largely complete. It can generate dossiers for a large subset of C++ including all annotations described above. The omissions are due mainly to the highly decomposed nature of the Roskind grammar (i.e., rare or obscure grammar productions have not been fleshed out). An initial version of the polymorphic load and store code is nearing completion (for a single platform). The interface to the persistent store has been defined and two distinct stores have been implemented. The first uses a version of DBM supporting transaction semantics. The other converts objects to a serial byte stream for use across interprocess communication channels. We plan to add an interface to the EXODUS storage manager[4] shortly.

Although the design described here is quite general there are a number of limitations in the current system. Most important, we do not support pointers to the interior of objects (although the load store hooks allow crude handling of some cases). We also do not support unions in the current system. Only two styles of dynamically sized data members are supported although many others can be envisioned. We are dissatisfied with the treatment of static data members mainly due to the uncertain semantics of persistent, shared members.

In terms of portability and simplicity of the solution there are several shortcomings. Of these, the most important is the requirement that the application programmer alter class definitions to include a reconstructor, load/store hooks (optional), and the dossier accessor function (optional). Another problem is the possibility that the byte order and floating point format must be explicitly indicated in the dossier making it non-portable.

## 10 Future work

The most important features currently unavailable in our system are heterogeneity and class evolution. To provide a universal and stable POS these are fundamental requirements. The design of these features is largely complete and an initial implementation should be completed soon. We hope to support both a simple static evolution algorithm and the dynamic one described in Section 8. We are also investigating the ability to lazily load individual nodes of the object graph. Given our current implementation constraints this will probably require complete object encapsulation. In addition, dynamically loading class definitions in the form of dossiers and member functions is possible through the use of our object/meta-object server[18].

A portable, comprehensive dossier facility has applications in a variety of areas. Two applications related to our research are inter-language object transmission[16] and dynamic reconfiguration of software systems[5].

## 11 Related Work

Persistent objects has been an area of intense research over the last few years and there are a large number of approaches. Table 1 provides a brief summary of some of these systems. A more in depth discussion of four representative systems is provided in the full paper.

System	Description Language	Dossiers	Preprocessor	Invocation	Implementation	Graph Traversal
Arjuna [8, 28]	Restricted C++	no	yes	special base class	rpc	no
Avalon [10]	Augmented C++	no	yes	special base class, <b>stable</b> keyword	rpc w/ transactions	inline code in r/w
C** [3, 17]	none	yes, not user visible	yes	object register method	vm and pointer swizzling	yes
E [19, 21, 20]	none	no	modified g++	parallel class hierarchy	vm and pointer swizzling	n/a
EC++ [24]	Restricted C++	no	yes	named object, special base class	rpc	inline code in r/w
NIHCL [12]	none	no	no	special base class, r/w functions	ASCII files	inline code in r/w
O++ [11, 7, 1]	Augmented C++	yes	compiler	overloaded new, special base class	tagged byte stream	yes
ObjectStore [14]	DB schema	yes	yes	overloaded new	vm and pointer swizzling	no
OBST [6, 30, 23]	Augmented C++	no	yes	create in container object	copied on container commit	names as roots
SOS [26, 25, 27]	Augmented C++	no	modified g++	special base class, overloaded new	object fault on special pointer class	special pointer class
Texas [29, 31]	a.out	yes, packed in 2 tables	tdesc	overloaded new	vm and pointer swizzling	n/a?
Utah Dossiers	Augmented C++	yes, persistent objects	Roskind grammar-based (fufu)	r/w functions	distributed rpc w/ tagged byte stream	dossier driven

Table 1: Summary of persistent objects systems and their approach.

## 12 Conclusions

By using dossiers as the foundation for a persistent object store we have built a flexible, portable storage facility capable of supporting class evolution and platform heterogeneity. The dossier generator requires minimal alteration of class descriptions and can be used where implementation source code is not available. Furthermore, the ability to manipulate objects polymorphically allows us to serialize arbitrary object graphs and restore them providing the basis for inter-process object transmission and RPC stub generation. A prototype of the dossier generator, polymorphic I/O code, and object store are nearing completion and initial experiments are encouraging.

## References

- [1] Rakesh Agrawal, Shaul Dar, and Narain H. Gehani. The o++ database programming language: Implementation and experience. In *Proceedings of the IEEE 9th International Conference on Data Engineering*. IEEE Computer Press, 1993.
- [2] Alpha\_1 Project. Integrated computer aided design and manufacturing: An overview of Alpha\_1. Technical report, University of Utah, Dept. of Computer Science, March 5, 1992.
- [3] Vinny Cahill, Chris Horn, Andre Kramer, Maurice Martin, and Gradimir Starovic. C\*\* and eiffel\*\*: languages for distribution and persistence. In *Proceedings of the 1990 OSF Microkernel Applications Workshop*, Grenoble, France, 1990.
- [4] Michael J. Carey, David J. DeWitt, Joel E. Richardson, and Eugene J. Shekita. Storage management for objects in EXODUS. In Won Kim and Frederick H. Lochovsky, editors, *Object-Oriented Concepts, Databases, and Applications*, pages 341–369. Addison-Wesley, 1989.
- [5] John B. Carter, Bryan Ford, Mike Hibler, Ravindra Kuramkote, Jeffrey Law, Jay Lepreau, Douglas B. Orr, Leigh Stoller, and Mark Swanson. FLEX: A tool for building efficient and flexible systems. In *Proc. Fourth Workshop on Workstation Operating Systems*, October 1993.
- [6] Eduardo Casais, Michael Ranft, Bernhard Schiefer, Dietmar Theobald, and Walter Zimmer. OBST – an overview. Technical report, Forschungszentrum Informatik (FZI), D-76131 Karlsruhe, Germany, 1993.
- [7] S. Dar, N. H. Gehani, and H. V. Jagadish. CQL++: A SQL for a c++ based object-oriented DBMS. In A. Pirotte, C. Delobel, and G. Gottlob, editors, *Advances in Database Technology – EDBT '92: Proceedings of the 3rd International Conference on Extending Database Technology*, Vienna, Austria, March, 1992, 1992. Springer-Verlag.
- [8] G.N. Dixon, G.D. Parrington, S.K. Shrivastava, and S.M. Wheeler. The treatment of persistent objects in Arjuna. In Stephen Cook, editor, *Proceedings of the 1989 European Conference on Object-Oriented Programming*, pages 169–189, University of Nottingham, July 10-14, 1989. Cambridge University Press.
- [9] Margaret A. Ellis and Bjarne Stroustrup. *The Annotated C++ Reference Manual*. Addison-Wesley, Reading, MA, 1990.
- [10] Jeffrey L. Eppinger, Lily B. Mummert, and Alfred Z. Spector, editors. *Camelot and Avalon: A Distributed Transaction Facility*. Data Management Systems. Morgan Kaufmann Publishers, Menlo Park, CA, 1991.
- [11] N. H. Gehani. OdeFS: A file system interface to an object-oriented database. Technical report, AT&T Bell Laboratories, Murray Hill, New Jersey 07974, 1989.

- [12] Keith E. Gorlen, Sanford M. Orlow, and Perry S. Plexico. *Data Abstraction and Object-Oriented Programming in C++*. John Wiley & Sons, 1990.
- [13] John A. Interrante and Mark A. Linton. Runtime access to type information in C++. In *USENIX Proceedings C++ Conference*, pages 233–240. USENIX Association, 1990.
- [14] Charles Lamb, Gordon Landis, Jack Orenstein, and Dan Weinreb. The ObjectStore database system. *Communications of the ACM*, 34(10):50–63, October 1991.
- [15] Robert W. Mecklenburg. The specification for a binary file format for alpha\_1 models. Alpha\_1 technical report 88-6, University of Utah, 1988.
- [16] Robert W. Mecklenburg. *Towards a Language Independent Object System*. PhD thesis, University of Utah, Salt Lake City, Utah, June 1991.
- [17] Michael Mock, Reinhold Kroeger, and Vinny Cahill. Implementing atomic objects with the Relax transaction facility. *Computing Systems*, 5(3):259–304, Summer 1992.
- [18] Douglas B. Orr and Robert W. Mecklenburg. OMOS — an object server for program execution. In *Proc. International Workshop on Object Oriented Operating Systems*, pages 200–209, Paris, September 1992. IEEE Computer Society. Also available as technical report UUCS-92-033.
- [19] Joel E. Richardson and Michael J. Carey. Persistence in the E language: Issues and implementation. *Software-Practice and Experience*, 19(12):1115–1150, December 1989.
- [20] Joel E. Richardson and Michael J. Carey. Implementing persistence in E. In John Rosenberg and David Koch, editors, *Persistent Object Systems: Proceedings of the Third International Workshop, Workshops in Computing*, pages 175–199. Springer-Verlag, Newcastle, Australia, January 10-13, 1989, 1990.
- [21] Joel E. Richardson, Michael J. Carey, and Daniel T. Schuh. The design of the E programming language. Technical Report 814, Computer Science Department, University of Wisconsin, Madison, WI, February 1989.
- [22] Jim Roskind. A yacc-able c++ 2.1 grammar, and the resulting ambiguities. July 1991.
- [23] Bernhard Schiefer, Dietmar Theobald, and Jürgen Uhl. User’s guide: OBST release 3.3. Technical report, Forschungszentrum Informatik (FZI), D-76131 Karlsruhe, Germany, July 1993.
- [24] Manuel Sequeira and José Alves Marques. Can c++ be used for programming distributed and persistent objects? In *Proceedings 1991 International Workshop on Object Orientation in Operating Systems*, pages 173–176, Palo Alto, CA, October 17-18, 1991. IEEE Computer Society Press.
- [25] Marc Shapiro. Prototyping a distributed object-oriented operating system on Unix. In *Proceedings of the First USENIX/SERC Workshop on Experiences with Distributed and Multiprocessor Systems*, pages 311–331, Fort Lauderdale, FL, October 5-6, 1989. Usenix Association.
- [26] Marc Shapiro, Yvon Gourhant, Sabine Habert, Laurence Mosseri, Michel Ruffin, and Céline Valot. SOS: An object-oriented operating systems—assessment and perspectives. *Computing Systems*, 2(4):287–337, Fall 1989.
- [27] Marc Shapiro and Laurence Mosseri. A simple object storage system. In John Rosenberg and David Koch, editors, *Persistent Object Systems: Proceedings of the Third International Workshop, Workshops in Computing*, pages 272–276. Springer-Verlag, Newcastle, Australia, January 10-13, 1989, 1990.

- [28] Santosh K. Shrivastava et al. *The Arjuna System Programmer's Guide*. Arjuna Research Group, Computing Laboratory, University of Newcastle upon Tyne, UK, February 1992. Public Release 1.0.
- [29] Vivek Singhal, Sheetal V. Kakkad, and Paul R. Wilson. Texas: An efficient, portable persistent store. In *Proceedings of The Fifth International Workshop on Persistent Object Systems (POS-V)*, San Miniato, Italy, September, 1992, 1992.
- [30] Jürgen Uhl, Dietmar Theobald, Bernhard Schiefer, Michael Ranft, Walter Zimmer, and Jochen Alt. The object management system of STONE: OBST release 3.3. Technical report, Forschungszentrum Informatik (FZI), D-76131 Karlsruhe, Germany, July 1993.
- [31] Paul R. Wilson and Sheetal V. Kakkad. Pointer swizzling at page fault time: Efficiently and compatibly supporting huge address spaces on standard hardware. In *Proceedings of the Second International Workshop on Object Orientation in Operating Systems*, pages 364–377, Dourdan, France, September 24-25, 1992. IEEE Computer Society Press.