

Distributed EZ

Alvaro E. Campos and David R. Hanson
*Department of Computer Science, Princeton University,
Princeton, NJ 08544*

Research Report CS-TR-349-91
September 1991

Abstract

EZ is a system that integrates traditional operating systems and programming languages into a very high-level, persistent, string processing language. This paper describes the design and initial implementation of a distributed memory manager that distributes *EZ*'s virtual address space transparently among a network of homogeneous computers. The design adapts the techniques used in recent implementations of shared virtual memory for use in *EZ*'s persistent environment. Unlike most implementations of shared virtual memory, control information is distributed and migrates. This memory manager works in concert with a distributed mark-and-sweep garbage collector, which is also concurrent and real-time. This collector trades time for space and minimal disruption of mutators, which reduces communication costs.

Introduction

A system that integrates language and operating system concepts into a single system offers a different perspective on some time-honored features of traditional operating systems. Traditional file systems are an example; there is no reason why files cannot be regarded simply as persistent values bound to variables. Once this “traditional barrier” is breached, it is natural to view all objects — records, arrays, procedures — as persistent values, which leads to one language that supports both programming and manipulating the environment.

EZ is a language-based, exploratory programming environment that supports this view [10, 17]. It accomplishes this integration by encapsulating system services as features in a very high-level language, and by providing a large, persistent virtual address space.

Previous research on *EZ* focused on language *design* — finding suitable linguistic encapsulations for system services. This approach has worked well for some services, e.g., file systems and processes, but other services strain this approach. Tough examples include interactive devices and distributed systems.

This paper reports on the different, but complementary, approach used to accommodate distribution in *EZ*. Instead of language design, this approach focuses on language *implementation* — finding suitable techniques for distributing the *EZ* virtual address space transparently among a network of homogeneous computers.

The EZ System

EZ is a high-level string processing language with a persistent memory that permits it to double as a programming environment. Services provided by traditional operating systems are cast as *EZ* language features. Values exist indefinitely or until changed, so *EZ*'s strings and associative tables subsume traditional file and directory services. Associative tables are also used for procedure activations [10] and for threads, i.e., lightweight processes, and low-level synchronization [17].

References [10] and [17] describe the syntax and semantics of *EZ* in detail. Briefly, *EZ* is a high-level string-processing language derived from Icon [13] and its predecessors. It shares many of their attributes, such as run-time flexibility, typed values and untyped variables, heterogeneous structures,

automatic type conversions, and mechanisms for runtime scope control. *EZ*'s built-in types includes numerics, strings, procedures, and associative tables. *EZ* has the usual control constructs, which are driven by the absence or presence of values. Everything in *EZ*, including program code, resides in a single, persistent, 32-bit virtual address space.

EZ is not the only system that offers persistence. APL systems have always operated in a persistent workspace, and some Smalltalk and Lisp systems offer similar facilities [22, 25, 32, 34]. Others have added persistent data types and procedures to Pascal-like languages [3, 4], some languages have mechanisms for transmitting arbitrary values between programs and hence long-term storage [18], and most object-oriented languages have facilities for saving some objects on disk for later retrieval [23, 28]. Persistence is also an important aspect of some programming environments for traditional languages [35] and for maintaining programming environment databases [33]. The key difference between these systems and *EZ* is that most offer persistence as a separate facility and usually restrict it to data. Persistence pervades *EZ* and applies to both data and active objects, like procedure activations and processes.

As described in Reference [17], *EZ* has an interpreter-based implementation similar in detail to other very high-level languages [12]. The virtual address space resides on disk and is managed by a virtual memory manager that caches pages in memory. Currently, all memory management is done by software. The interface between the interpreter and memory manager consists of two functions: `GetPage(a, mode)` returns a “handle” to the page that encompasses address *a*, and `PutPage(h, dirty)` releases the page given by handle *h*. *mode* indicates access mode and is either `read` or `write`, and *dirty* indicates whether or not a writable page was actually modified. Handles are simply the memory address of the in-cache copy of the page. Flushing the cache saves the system state.

Unlike the clients of memory managers in operating systems, the *EZ* interpreter is a “trusted” client and is therefore given complete ownership of pages requested via `GetPage`. Other interpreter threads can request the same page. As in classical readers/writers applications, multiple `read` requests are granted, but `write` requests are given exclusive access. The interpreter itself is written to avoid deadlock, but it is possible to write *EZ* processes that deadlock just as it is possible to write programs with threads in Mach [6] that deadlock, for example.

Distribution

The *EZ* address space is distributed transparently among a set of homogeneous processors by replacing the memory manager by one that is based on the recent implementations of shared virtual memory [20, 21, 24, 29]. The shared virtual memory manager attempts to collect recently accessed pages at the processor that accesses them. In this design, identical copies of the new memory manager run at each processor, and the entire virtual address space is distributed among the secondary storage devices of the individual processors. Each manager caches the pages that its interpreter clients are using, and these pages are accessed via `GetPage` as described above. The managers hide all distribution details. Managers replicate pages to permit multiple readers, but permit only one writer in order to maintain coherence.

Coherence

Since `GetPage` yields ownership of a page to the client interpreter, the managers cannot use invalidation techniques [20] to maintain coherence. Page ownership migrates to the manager that needs to grant write access to the page. At any time, each page is owned by exactly one manager, and that owner is responsible for maintaining the disk copy of the page. Consequently, pages migrate among local disks.

Each cached page is known to a manager as *invalid*, *shared*, or *unique*. Invalid pages are those owned by other managers for which the cached copy is known to be out of date. A new copy must be fetched from its owner the next time access is requested. Shared pages are those for which the cached copy might be up to date and other managers might also have shared copies. Read requests for shared pages can be granted after the owner confirms the validity of the local copy. Unique pages are those that are owned by the local manager and for which there are no valid copies elsewhere. Read and write requests for these pages can be granted without communicating with other managers.

As suggested above, a manager maintains a page cache and an ownership table, which lists all of the pages owned by the manager and their disk addresses. For each cached page, the cache entry holds a pointer to the local copy of the page, the page status (invalid, shared, or unique), a reference count, an ownership hint, a valid copies hint, a dirty bit, and a queue of pending access requests. The last three entries apply only to locally owned pages.

For remote pages (i.e., pages owned by another manager), the reference count is the number of read accesses that have been granted in response to `GetPages` from local interpreter threads. For owned (and hence local) pages, this count is the number of local read and write accesses that have been granted plus the number of remote managers that have granted access to their copies of the page (which they obtain on the first access request). `PutPages` decrement the count. For remote pages, a message is sent to the owner when the count reaches 0, and the owner then decrements its counter.

For owned pages, write requests are queued until the count becomes 0. Once granted, the count becomes 1. For remote pages, write requests cause the local manager to ask the page's owner to transfer ownership to the local manager, which occurs when the owner's count becomes 0. Once transferred, the write request is granted as above.

The ownership hint identifies owned pages and gives the probable identities of their owners. The owner's identity accompanies the copies of remote pages when they are fetched. Responses always include the owner's identity, which updates this field, if necessary.

The valid copies hint is a bitmap that identifies other managers to which valid copies of the page have been sent. It is used to avoid sending copies of pages unnecessarily.

Pages with a 0 reference count and no pending requests may appear on one of four LRU lists, which are consulted for page replacement. When the manager must replace a cache entry, it uses the first entry on the list of *invalid* pages. If this list is empty, it uses the first entry on list of remote *shared* pages. Doing so obligates the local manager to fetch a new copy of the replaced page, if it is accessed again.

If both of these lists are empty, the first entry on the list of *unique* pages is used. Selecting a page from this list obligates the manager to reread it, if it is accessed again, but there is no network cost associated with this choice.

If all else fails, the first entry on the list of owned *shared* pages is used. There are valid copies of these pages elsewhere, but the manager must "forget" about these copies because it is about to reuse the valid copies hint in the cache entry. Consequently, subsequent read requests from remote managers will cause the page to be resent, perhaps unnecessarily.

Communication

Communication between managers is based on reliable, order-preserving datagrams. Request messages have a *type*, which identifies the request,

a *sender*, which identifies the initiating manager and its processor, and an *address*, which is the virtual address of interest in the request. Acknowledgment messages have a *sender*, which identifies the owner of the page, and an optional *page*, which is a copy of the page itself.

The sender helps managers differentiate between local and remote requests. The addresses in remote requests always specify pages; local requests mirror the semantics of `GetPage` and can specify any virtual address and the page containing that address is returned.

Calls to `GetPage` and `PutPage` generate four message types. `GetPage` generates *read* and *write* requests. `PutPage` generates *clean* requests, which indicate that the client is finished with the page and has not modified it, and *dirty* requests, which indicate that the client is finished with the page and modified it.

Most requests are for pages that the local manager owns, and handling these is straightforward. Requests for pages that the local manager does *not* own require communication with other managers, specifically the owner of the requested page. There are seven message types involved in this communication, which are refinements of the four types described above.

`GetPage` for *read* access causes the local manager to send a *copy* request to the page's owner who grants the request by sending a copy of the page. If the local manager already has a copy of the page that it believes is valid, a *read* request is sent to the owner. The owner simply grants access, but includes a copy of the page if it is possible that the requester's copy is invalid. When a read-only page is returned via `PutPage`, the local manager decrements its reference count on the page and, when the count reaches 0, sends the owner a *decrement* request, which is the remote equivalent of *clean*.

Calls to `GetPage` that request *write* access causes page ownership to transfer to the requester. The local manager sends the owner a *write* or *fetch* request depending on whether or not it holds a valid copy of the page. The owner grants the request, including a copy of the page, if necessary, and marks its copy of the page as *invalid* since it is no longer the owner and its copy of the page is now obsolete.

If a requesting manager does not know the owner of the desired page (e.g., because it does not hold a valid copy of the page), it broadcasts equivalents of the message types described above to all managers. Only the actual owner responds to these messages, which identifies the owner for subsequent requests and reduces broadcast traffic. Messages directed to a manager that is no longer the owner of the desired page are turned into broadcast messages

by the recipient or simply forwarded if the recipient has an ownership hint for the page.

Implementation

The implementation of the memory manager consists of about 900 lines of ANSI C, not including the thread and communication packages. The interpreter and runtime system consist of another 5,000 lines.

Each cached page is served by a thread, which processes requests from a per-page FIFO service queue. When the queue becomes empty, the thread terminates. Requests that cannot be satisfied immediately because the page is busy are queued and serviced when previous requests complete. Using one thread per page serializes requests and simplifies programming.

Other threads accept requests from remote managers and append them to the service queue of the appropriate page. These threads also initiate per-page threads as necessary.

Garbage Collection

Earlier versions of *EZ* used an off-line garbage collector to reclaim inaccessible pages in the disk representation of the virtual address space. This approach is fine for a prototype and, as shown by the Oberon system [37], perhaps adequate for a non-distributed, "single-user" *EZ* system that is subject to frequent idle periods, but an off-line approach is unsuitable for a distributed system.

Distributed *EZ* will use a distributed garbage collector that works in concert with the shared virtual memory manager described above. It is a distributed mark-and-sweep collector [5, 7, 27], and it is concurrent and real-time. Technically, algorithms based on reference counting [11, 19] are more efficient, but require additional data for every pointer that might refer to a page on another processor, or additional synchronization between subsets of the processors at each reference. Besides, these algorithms cannot handle cycles, which makes them unsuitable for *EZ* where cycles abound.

Likewise, copying collectors are also more efficient and can be made both concurrent and real-time [1]. Most designed for distributed address spaces are not concurrent nor real-time, and some require special hardware to be efficient [14] or impose restrictions on inter-processor pointers, such as double indirection [26]. More importantly, objects in *EZ*'s virtual address

space cannot move; this restriction greatly simplifies the implementation of the persistent address space.

For systems with large, persistent address spaces, efficiency is less important than concurrency. Indeed, all that is required is that the collector replenish the supply of free pages fast enough so that applications rarely have to wait to allocate a new page, and that it eventually collects all inaccessible pages.

An identical copy of the collector runs forever on each processor (technically, there is a collector for each memory manager). As usual, a collector marks all pages that hold accessible objects starting from a few system "root" objects. It also marks pages referenced from within objects on marked pages, which may cause some inaccessible pages to be marked. Indeed, the collector is conservative: it marks a superset of the accessible pages and collects only a subset of the inaccessible ones [7]. It repeats the collection continuously, so it eventually reclaims all pages that are not referenced by accessible pages.

Each collector processes only the pages owned by its cooperating memory manager. After marking the accessible owned pages, a collector exchanges information about inter-processor references with the collectors on the other processors. This information feeds another marking cycle that expands the local collector's set of accessible pages. This activity continues until no collector can expand its accessible page set.

Concurrent collectors need the cooperation of mutators when updating references. In particular, when a reference is updated, either the old target or the new target must be marked atomically. Failure to do so may cause the collector to reclaim either the old or new target erroneously. Most collectors mark the new target to avoid hanging on to the page holding the old target unnecessarily. *EZ's* collector, however, marks the old target for two reasons. First, marking the new target requires direct mutator assistance. Second, marking the old target permits the collector to use virtual memory hardware to mark pages referenced by a page before it is modified [2]. At the beginning of a collection, all owned pages are set to read only. The first write to a page causes a page fault, and pages referenced within the faulted page are marked before the fault handler approves write access to the page. Marking the referents of a page *before* it is updated is equivalent to marking the old targets of updates. Similar comments apply to calls to `GetPage` for write access, but virtual memory hardware is unnecessary.

At any time during marking, each page is marked with either white, grey, or black. Stop-and-collect, uniprocessor collectors that are not concurrent

use only two colors; the third, *grey*, labels pages that have been marked but not yet traversed for pointers.

At the beginning of each collection, locally owned pages have the same color, say, white. After marking, pages will be colored either white or black, and white pages can be reclaimed. As suggested above, some black pages may really be inaccessible, but they will be reclaimed during a subsequent collection. Pseudo-code for the collection algorithm is shown in Figure 1. Initially, *retain* and *gather* are black and white, respectively. The colors reverse roles in subsequent collections.

Locally free pages are colored black so that the memory manager can allocate them without help from the collector. Owned pages are set to *readonly* so that their referents will be marked if they are modified, as described above. The referents of the local roots are colored by *Shade*, which colors owned white pages grey or adds remote pages to *m*, which is empty initially.

Marking then begins to cycle. In each cycle, locally owned grey pages are scanned. As shown in Figure 1, scanning a page shades its referents, colors the page black, and unprotects it. Scanning a page may yield another grey page, but this activity ends eventually. Once all grey owned pages are scanned, all pages reachable from local roots are colored black, and *m* is the set of all remote pages that should have been colored grey.

A subset of *m* is broadcast to other processors. This subset is the set of pages that have not been announced by previous broadcasts. *cycle* records the size of this subset, and *m* is added to *M*, which accumulates remote pages announced by any processor. The collector then consumes similar messages from the other collectors, accumulates their sizes in *cycle*, and shades the pages mentioned in the messages. This shading colors owned white pages grey or adds remote pages to *M*. These messages serve not only to communicate the remote page references between collectors, but also to synchronize them.

As collection progresses, each collector's *M* becomes larger until *m - M* becomes empty, i.e., until the successors of all grey pages everywhere have been colored black. At that point, all owned white pages are added to the set of free pages, the roles of white and black are reversed and collection begins anew.

Space efficiency and minimal disruption of mutators are more important than time efficiency of the algorithm itself. There are reasonably efficient representations for all of the data structures used in the algorithm. Marks are kept in a private bitmap, 2 bits per page. Page sets for *k*, *m*, and *M*

```

retain ← black
gather ← white
do forever
   $M \leftarrow m \leftarrow \emptyset$ 
  for every  $p \in \text{free pages}$  do  $\text{Color}(p) \leftarrow \text{retain}$ 
  for every  $p \in \text{owned}$  do  $\text{Access}(p) \leftarrow \text{readonly}$ 
  for every reference  $r$  in the local roots do  $\text{Shade}(\text{Page}(r), m)$ 
  do
    while there is a  $p \in \text{owned} \wedge \text{Color}(p) = \text{grey}$  do  $\text{Scan}(p, m)$ 
     $\text{cycle} \leftarrow |m - M|$ 
    broadcast  $m - M$ 
     $M \leftarrow M \cup m$ 
    for every other processor  $P$  do
      receive message  $k$  from  $P$ 
       $\text{cycle} \leftarrow \text{cycle} + |k|$ 
      for every  $p \in k$  do  $\text{Shade}(p, M)$ 
     $m \leftarrow \emptyset$ 
  while  $\text{cycle} > 0$ 
  for every  $p \in \text{owned}$  do
    if  $\text{Color}(p) = \text{gather}$  then  $\text{free pages} \leftarrow \text{free pages} \cup \{p\}$ 
   $\text{gather}, \text{retain} \leftarrow \text{retain}, \text{gather}$ 

Shade( $p, s$ ):
  if  $p \in \text{owned}$  then
    if  $\text{Color}(p) = \text{gather}$  then  $\text{Color}(p) \leftarrow \text{grey}$ 
  else  $s \leftarrow s \cup \{p\}$ 

Scan( $p, s$ ):
  for every reference  $r$  in  $p$  do  $\text{Shade}(\text{Page}(r), s)$ 
   $\text{Color}(p) \leftarrow \text{retain}$ 
   $\text{Access}(p) \leftarrow \text{read/write}$ 

```

Figure 1: Garbage Collection Algorithm.

are represented as lists of arrays of page numbers or page ranges, and most page ranges can fit in 32 bits as a 22-bit page and a 10-bit spread or as two 16-bit page numbers. The memory manager already maintains *owned* as a sorted list of page ranges, and *Shade* threads a list through these entries to maintain a set of owned *grey* pages. Finally, *free pages* is a list of available pages represented by a list of arrays of page numbers stored in disk blocks as in UNIX [31]. It is accessed as a LIFO list, so the first disk block is almost always in the memory manager's cache.

Discussion

Implementation of the garbage collector and the communications layer and integrating these components into the *EZ* system are underway. These changes also necessitated basic changes in the *EZ* virtual machine in order to insure atomicity and consistency.

Previously, many primitives that accessed shared memory (as opposed to per-thread memory) accepted pointers into the cache as operands because the previous version of *GetPage* was atomic and uninterruptible. The distributed memory manager forced these kinds of primitives to be decomposed and re-cast in terms of three operations on associative tables: membership testing, insertion, and deletion. These operations are atomic with respect to the tables on which they operate, but can be interrupted by other, unrelated operations.

Performance measurements will undoubtedly induce modifications to the design and to the current implementation. For example, even though relatively few shared pages are actually modified, *EZ's* adherence to strict consistency may lead to thrashing, which might be attacked by selective use of release consistency within the interpreter, or by instituting a minimum ownership time [9], which would give owners time to complete several modification operations before a page changed owner because of a write request.

Other than the interface via *GetPage* and *PutPage*, there is little in the virtual memory system that is specific to *EZ*. These techniques can be applied to other distributed persistent languages and environments.

EZ's 32-bit virtual address space is too small, especially in light of the impending availability of 64-bit processors. The techniques described in this paper can accommodate such large address spaces, but other techniques might have advantages. Pointer "swizzling" [36] could be used for a non-distributed *EZ*, but it is unclear how to adapt this technique to distributed

systems.

Another alternative design under consideration accommodates multiple virtual address spaces anywhere in a network and achieves distribution by inter-address space references. This approach can be viewed as the complement of the distributed memory manager approach. Here, the original memory manager remains nearly untouched, but the language, interpreter and runtime system are modified.

References to other address spaces are made through inter-address space pointers. These pointers are functionally equivalent to capabilities used in some distributed systems [30]. They contain an address space identifier and an address within that address space. Capability-like rights could be added to restrict the set of legal operations.

The attraction of this approach is that address spaces could be encapsulated as tables much like strings encapsulate files. The cost, however, is that the interpreter must take special action in order to access these tables. Efficient implementation techniques for this kind of dereferencing have been used in other high-level language systems [12, 16], in heterogeneous systems [8], and for implementing implicit synchronization [15]. Pages would not migrate in this approach, which simplifies memory management and increases reliability. Remote dereferencing translates into essentially remote procedure calls to `GetPage` and `PutPage`.

Reliability and fault tolerance for distributed persistent systems like *EZ* remains an important area for future work. Currently, distributed *EZ* uses timeouts to recover from network and machine failures. These timeouts cause the memory manager to terminate the requesting *EZ* process. This somewhat unsatisfying approach works for user-level processes, but is unacceptable for system-level processes, like the garbage collector, and other mechanisms are under investigation.

References

- [1] A. W. Appel, J. R. Ellis, and K. Li. Real-time concurrent collection on stock multiprocessors. *Proceedings of the '88 SIGPLAN Conference on Programming Language Design and Implementation, SIGPLAN Notices*, 23(7):11–20, July 1988.
- [2] A. W. Appel and K. Li. Virtual memory primitives for user programs. In *Proceedings of the Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 96–107, Santa Clara, CA, Apr. 1991.

- [3] M. P. Atkinson, K. Chisholm, P. Cockshott, and R. Marshall. Algorithms for a persistent heap. *Software—Practice & Experience*, 13(3):259–271, Mar. 1983.
- [4] M. P. Atkinson and R. Morrison. Procedures as persistent data objects. *ACM Transactions on Programming Languages and Systems*, 7(4):539–559, Oct. 1985.
- [5] L. Augusteijn. Garbage collection in a distributed environment. In J. W. de Bakker, A. J. Nijman, and P. C. Treleaven, editors, *PARLE, Parallel Architectures and Languages Europe, Volume II: Parallel Languages (Lecture Notes in Computer Science 259)*. Springer-Verlag, Berlin, 1987.
- [6] E. C. Cooper and R. P. Draves. C threads. Technical Report CMU-CS-88-154, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, June 1988.
- [7] E. W. Dijkstra, L. Lamport, A. J. Martin, C. S. Scholten, and E. F. M. Steffens. On-the-fly garbage collection: An exercise in cooperation. *Communications of the ACM*, 21(11):966–975, Nov. 1978.
- [8] R. B. Esick IV. *The Cross-Architecture Procedure Call*. PhD thesis, The University of Illinois at Urbana-Champaign, Urbana, IL, 1987.
- [9] B. Fleisch and G. J. Popek. Mirage: A coherent distributed shared memory design. In *Symposium on Operating System Principles*, pages 211–223, Oct. 1989.
- [10] C. W. Fraser and D. R. Hanson. High-level language facilities for low-level services. In *Conference Record of the ACM Symposium on Principles of Programming Languages*, pages 217–224, New Orleans, LA, Jan. 1985.
- [11] B. Goldberg. Generational reference counting: A reduced-communication distributed storage reclamation scheme. *Proceedings SIGPLAN '89 Conference on Programming Language Design and Implementation, SIGPLAN Notices*, 24(7):313–321, July 1989.
- [12] R. E. Griswold and M. T. Griswold. *The Implementation of the Icon Programming Language*. Princeton University Press, Princeton, NJ, 1986.
- [13] R. E. Griswold and M. T. Griswold. *The Icon Programming Language*. Prentice Hall, Englewood Cliffs, NJ, second edition, 1990.
- [14] R. H. Halstead, Jr. Implementation of Multilisp: Lisp on a multiprocessor. In *Conference Record of the 1984 ACM Symposium on LISP and Functional Programming*, pages 9–17, Austin, TX, Aug. 1984.
- [15] R. H. Halstead, Jr. Multilisp: A language for concurrent symbolic computation. *ACM Transactions on Programming Languages and Systems*, 7(4):501–559, Oct. 1986.

- [16] D. R. Hanson. Variable associations in SNOBOL4. *Software—Practice & Experience*, 6(2):245–254, Apr. 1976.
- [17] D. R. Hanson and M. Kobayashi. EZ processes. In *Proceedings of the International Conference on Computer Languages*, pages 90–97, New Orleans, LA, Mar. 1990.
- [18] M. Herlihy and B. H. Liskov. A value transmission method for abstract data types. *ACM Transactions on Programming Languages and Systems*, 4(4):527–551, Oct. 1982.
- [19] C.-W. Lermen and D. Maurer. A protocol for distributed reference counting. In *Proceedings of the 1986 ACM Symposium on LISP and Functional Programming*, pages 343–350, Cambridge, MA, Aug. 1986.
- [20] K. Li and P. Hudak. Memory coherence in shared virtual memory systems. *ACM Transactions on Computer Systems*, 7(4):321–359, Nov. 1989.
- [21] K. Li and R. Schaefer. A hypercube shared virtual memory system. In *Proceedings of the International Conference on Parallel Processing*, pages 125–132, University Park, PA, 1989.
- [22] C. Low. A shared, persistent object store. In S. Gjessing and K. Nygaard, editors, *ECOOP '88: European Conference on Object-Oriented Programming (LNCS 322)*, pages 390–410, Berlin, Aug. 1988. Springer-Verlag.
- [23] B. Meyer. *Object-Oriented Software Construction*. Prentice Hall International, London, 1988.
- [24] B. Nitzberg and V. Lo. Distributed shared memory: A survey of issues and algorithms. *IEEE Computer*, 24(8):52–60, Aug. 1991.
- [25] A. Paepcke. PCLOS: A flexible implementation of CLOS persistence. In S. Gjessing and K. Nygaard, editors, *ECOOP '88: European Conference on Object-Oriented Programming (LNCS 322)*, pages 374–389, Berlin, Aug. 1988. Springer-Verlag.
- [26] M. Rudalics. Distributed copying garbage collection. In *Proceedings of the 1986 ACM Symposium on LISP and Functional Programming*, pages 364–372, Cambridge, MA, Aug. 1986.
- [27] M. Schelvis and E. Bledsoe. The implementation of a distributed Smalltalk. In S. Gjessing and K. Nygaard, editors, *ECOOP '88: European Conference on Object-Oriented Programming (LNCS 322)*, pages 374–389, Berlin, Aug. 1988. Springer-Verlag.
- [28] Stepstone Corp., Sandy Hook, CT. *Objective-C Compiler Version 4.0: User Reference Manual*, Sept. 1988.
- [29] M. Stumm and S. Zhou. Algorithms implementing distributed shared memory. *IEEE Computer*, 23(5):54–64, May 1990.

- [30] A. S. Tanenbaum, R. van Renesse, H. van Staveren, G. J. Sharp, S. J. Mullender, J. Jansen, and G. van Rossum. Experiences with the Amoeba distributed operating system. *Communications of the ACM*, 33(12):46–63, Dec. 1990.
- [31] K. Thompson. UNIX implementation. *Bell System Technical Journal*, 57(6):1931–1946, July 1978.
- [32] S. R. Vegdahl. Moving structures between Smalltalk images. *OOPSLA '86 Conference Proceedings, SIGPLAN Notices*, 21(11):466–471, Nov. 1986.
- [33] D. Weibe. A distributed repository for immutable persistent objects. *OOPSLA '86 Conference Proceedings, SIGPLAN Notices*, 21(11):453–465, Nov. 1986.
- [34] D. S. Wile, N. Goldman, and D. G. Allard. Maintaining object persistence in the Common Lisp framework. In *Proceedings of the Second International Workshop on Persistent Object Systems*, pages 382–405, Appin, Scotland, Aug. 1987.
- [35] J. C. Wileden, A. L. Wolf, C. D. Fisher, and P. L. Tarr. PGRAPHITE: An experiment in persistent typed object management. *Proceedings of SIGSOFT '88: Third Symposium on Software Development Environments, SIGPLAN Notices*, 24(2):130–142, Feb. 1989.
- [36] P. R. Wilson. Pointer swizzling at page fault time: Efficiently supporting huge address spaces on standard hardware. *Computer Architecture News*, 19(4):6–13, June 1991.
- [37] N. Wirth and J. Gutknecht. The Oberon system. *Software—Practice & Experience*, 19(9):657–693, Sept. 1989.