

Orthogonally Persistent Object Systems

Malcolm Atkinson[¥] and Ronald Morrison[†]

[¥]Department of Computer Science, University of Glasgow,
Glasgow G12 8QQ, Scotland
mpa@dcs.glasgow.ac.uk

[†]School of Mathematical and Computational Sciences,
University of St Andrews, North Haugh, St Andrews KY16 9SS, Scotland
ron@dcs.st-andrews.ac.uk

Keywords: orthogonal persistence, persistent programming languages, database programming languages, persistent application systems

This paper should be referenced as:

Atkinson, M.P. & Morrison, R. "Orthogonally Persistent Object Systems". VLDB Journal 4, 3 (1995) pp 319-401.

ABSTRACT

Persistent Application Systems (PASs) are of increasing social and economic importance. They have the potential to be long-lived, concurrently accessed and consist of large bodies of data and programs. Typical examples of PASs are CAD/CAM systems, office automation, CASE tools, software engineering environments and patient-care support systems in hospitals. Orthogonally persistent object systems are intended to provide improved support for the design, construction, maintenance and operation of PASs. The persistence abstraction allows the creation and manipulation of data in a manner that is independent of its lifetime thereby integrating the database view of information with the programming language view. This yields a number of advantages in terms of orthogonal design and programmer productivity which are beneficial for PASs.

Design principles have been proposed for persistent systems. By following these principles, languages that provide persistence as a basic abstraction have been developed. In this paper the motivation for orthogonal persistence is reviewed along with the above mentioned design principles. The concepts for integrating programming languages and databases through the persistence abstraction and their benefits are given. The technology to support persistence, the achievements and future directions of persistence research are then discussed.

1 Introduction

This paper presents a broad review of current research and achievements in orthogonally persistent object systems. It provides an advanced tutorial for those commencing research or advanced study in persistence, databases or database programming languages and a survey of the persistent language community's contribution to this growing research area. A particular goal is to summarise the existing achievements and to identify the current research issues.

Orthogonally persistent object systems support a *uniform* treatment of objects irrespective of their types by allowing values of *all* types to have whatever longevity is required. The motive for establishing this uniform treatment is presented first.

Subsequent sections of the paper present definitions, integration concepts, technology to support persistence and achievements. The penultimate section suggests some of the crucial areas of further research. Table 1 identifies the major topics covered.

Table 1: Location of Major Topics in this Paper

Topic	Section
Problems with Conventional Technology	1.1&1.2
Principles of Persistence & Definition of Orthogonal Persistence	2
Benefits of Orthogonal Persistence	2.3
Integration of Databases & Programming Languages	3
Types & Data Models	3.1
Persistent Type Systems	3.1.1
Binding Mechanisms	3.2
Concurrency and Transactions	3.3
Technology to Support Orthogonal Persistence	4
Implementation Architectures	4.1
Implementing Persistence by Reachability	4.2
Type-Safe Linguistic Reflection	4.3
Incremental Persistent Application System Construction	5.1
Hyper-Programming	5.2
Persistent Software Engineering	5.3&5.4
Research Directions	6
Extensions of Persistence	6.1
Exploitation of Persistence	6.2
Delivering Persistence	6.3
Summary & Conclusions	7

1.1 Incoherence of Current Technology

The aim of persistent programming is to support the activity of applications' construction for long-lived, concurrently accessed and potentially large bodies of data and programs; referred to here as Persistent Applications Systems (PASs). They are designated persistent *application* systems because the *application* often outlives its individual components and even its implementation technology. Typical examples of such PASs are: CAD/CAM systems, office automation, CASE tools, software engineering environments [Teitelbaum & Reps, 1981, Akima & Ooi, 1989, Bott, 1989, Sommerville *et al.*, 1989, Thomas, 1989], integrated hospital administration and medical systems, large scientific databases and programs that analyse them, geographic information systems, environmental modelling systems, object-oriented databases [Bancilhon *et al.*, 1988, Bretl *et al.*, 1989] and process modelling systems [Bruynooghe *et al.*, 1991, Curtis *et al.*, 1992, Han & Welsh, 1993].

Those that commission the construction of a PAS expect it to be built at reasonable cost, to be delivered on time and to serve their organisation reliably for many years, accommodating change as requirements change and as technology advances. They are often disappointed since it frequently proves much harder and more expensive to build and maintain a PAS than was expected and its evolution is invariably problematic.

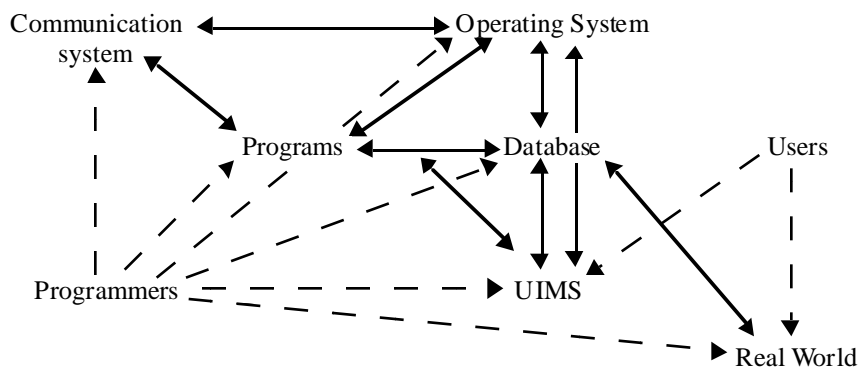
This may be illustrated with an example. In a health care management system in use in a hospital, the medical staff were alarmed when many of the fields in the medical records disappeared. Their local support staff diagnosed a fault in the database and restored it from an earlier state. The problem *persisted*, so they made an earlier restoration, losing more information. Still the problem persisted. The suppliers of the PAS were called in and eventually discovered that a file containing a font used by the user interface management system had been lost. Inconsistent behaviour by the supporting subsystems had led to a very expensive service interruption and the loss of information. If restoration had functioned identically for all the data, then the lost file would have been restored. As different data was treated differently, the local staff found the system incomprehensible and the implementors had a struggle to understand the cause of the failure. Similar, but sometimes more detailed inconsistencies, bedevil every stage of PAS construction and maintenance when it is constructed using traditional technology.

Currently the technology underlying PAS building (called “the support system” in this paper) relies on a number of disparate mechanisms and philosophical assumptions for support and efficient implementation [Atkinson *et al.*, 1983]. Among these are operating systems, communications systems (often bundled with the operating system but designed separately), database systems, user interface systems, command languages, editors, file systems, compilers, interpreters, linkage editors, binders, debuggers, DBMS — DDLs and DMLs, query languages, user interface management systems (UIMS), transaction managers, concurrency managers and machine types. Programmers have to cope with variations such as different naming, type and binding schemes in each of the components they use. Perhaps the hardest challenge these programmers face is coping with different recovery, concurrency and transactional behaviour.

The incoherence and complexity arising from utilising these many different and diverse application and system building mechanisms increases the cost both intellectually and mechanically of building even the simplest of systems. The complexity distracts the application builders from the task in hand forcing them to concentrate on mastering the multiplicity of programming systems rather than the application being developed. Perversely the plethora of disparate mechanisms is also costly in machine terms in that the code for interfacing them, their redundant duplication of facilities and their contention for resources cause execution overheads.

Figure 1 illustrates the interconnection of system components and users of the system. In moments of crisis and during peaks of activity we expect the most from our PASs. Unfortunately, it is precisely at such moments of stress that they malfunction due to the complexity of inter-dependencies between components.

Figure 1: The Complex Inter-dependence in Current Systems



In Figure 1, the solid lines represent data flow, translation and mappings, the dashed lines what people (users and programmers) are required to understand. The situation shown is ideal for the users, they are required only to understand the minimum. However, observations of real applications show that the users find they have to learn about aspects of the operating system and database. On the other hand, the situation is far from ideal for application programmers as they have to deal with *four* subsystems as well as the programs and the applications that are their real concern.

Individual inconsistencies between different underlying subsystems present software engineers working on a PAS with enough problems. However, each pair of subsystems presents these discontinuities and the composition of complexity is often multiplicative rather than additive. Recognition of the cost of such disharmonies has motivated the designers of orthogonally persistent object systems to propose and support computational models that allow large scale and long-lived computation without these unnecessary sources of complexity.

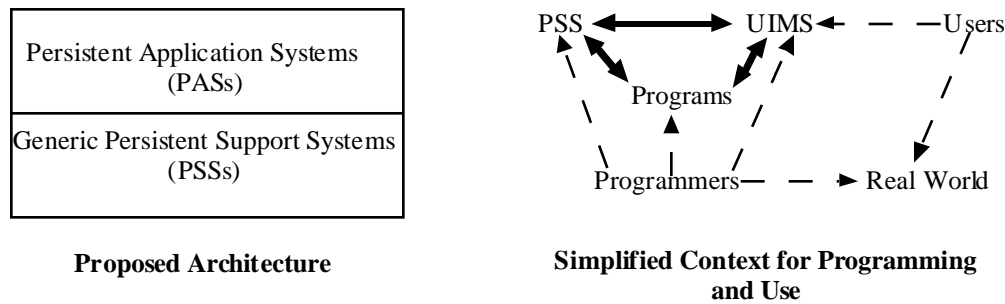
It transpires that, in many cases, the inconsistent treatments are not fundamental but accidental [Atkinson, 1978]. The various subsystems were built at different times

when the engineering trade-offs were different. In consequence, they provide virtually the same services, but inconsistently as they were designed and developed independently. By choosing to provide the total composition of services needed to support a PAS within one coherent design, the designers of orthogonally persistent object systems have eliminated the accidental disharmonies and exposed the few places where it is a difficult research issue to find the integrating concepts.

It is now timely to propose new models of computation, expressed both linguistically and architecturally, in response to the advent of new hardware technologies with different time, space and cost trade-offs and to the development of new application areas, such as office automation, CAD/CAM, CASE, etc. Persistent systems are one such proposal.

In summary, it is important to distinguish carefully between a PAS and the support system which enables it to operate and application programmers to build and maintain it. As the cost of building the support system can be amortised over many PASs and as the cost of application programming dominates the software industry, it is appropriate to invest in improved support systems to reduce PAS programming costs. Orthogonal persistence provides the design principles to guide this investment. The desired simple relationship between a PAS and a generic support system is shown in Figure 2.

Figure 2: Target Architecture to Support Persistent Application Systems



1.2 Analysis of the Causes of Incoherence

Some of the complexity in a PAS is intrinsic in that it emanates from the tasks which the PAS sets out to support. However, as argued above, a considerable part of this complexity is extraneous, an unfortunate artefact of the independent development of several essential supporting technologies. Before proposing integrated design principles which eliminate this extraneous complexity, its origins are analysed.

This analysis will focus on the relationship between programming languages and database systems for three reasons:

- it is typical of the class of problems that arise, e.g. comparing operating systems and programming languages would reveal similar extraneous differences;
- it is the domain that has received most attention; and
- it matches the interests of the expected audience of this paper.

The database and programming language communities have continued to research and develop products independently of one another despite having to provide many similar services. Each provides the means of naming values, each allocates space in which to store values, each provides a means of constraining those values and the operations on them, each provides modelling and description mechanisms, each provides mechanisms for extracting values from data structures, each provides concurrency, and so on.

The inconsistencies originate from different philosophies and circumstances in the two camps. The database community were faced with severe engineering problems. It is difficult to preserve large volumes of data reliably and to support many processes operating against that data efficiently. It is also difficult to accommodate the design habits of the system analyst and to allow virtually independent teams of programmers to develop parts of a PAS independently. Consequently, technical solutions dominate. Relations, query optimisation, serialisable transactions and views are examples of their success.

By contrast, programming language designers have typically aspired to help programmers be precise and to make programs understandable. They have sought design principles that lead to languages that have regular rules. These rules should match programmers' intuitions, be easy to define precisely and lead to programs that can be executed with reasonable efficiency.

Of course, there are exceptions in each camp. Relations have simple, formal and precise rules and there is an unfortunate number of programming languages which do not. The responses to design difficulties in each context is informative. In programming languages, essential requirements are often ignored, for example, many languages have had no predefined treatment of I/O and functional languages offer no mechanism for update. In databases, a solution is often designed independently and simply added, e.g. long-running transactions. These additions may not compose well with the existing constructs.

The separate development and consequent inconsistencies tend to perpetuate and grow. The intellectual and software investment in each camp militates against easy adoption of the other's ideas. The dichotomy between philosophies continues and may be heightened as they view each other through caricatures. To a programming language designer the database world looks like a mess of incomprehensible *ad hoc* design with little underlying philosophy. On the other hand, database designers are surprised that programming languages are so unhelpful with real problems such as bulk types, persistence, concurrency and transactions.

The challenge is to show that they need not remain behind separate barricades and that PAS builders can benefit from the resultant integrated support environment rather than suffer the task of bridging a perpetuated and possibly widening gulf.

1.3 Introduction to Persistence

Before introducing persistence, it may be important to point out one approach to the above incoherence that does not lead to a solution. That approach is to perpetuate the current underlying technologies, by gluing them together with enough glue-ware and hiding them behind sufficient "standard" interface veneers. The underlying differences in semantics ultimately show through as, for example, in failure semantics when the combined system is stressed. Since PAS builders wish to support near continuous availability or at least explain to users what is happening when service commitments are not met, they have to understand the underlying technology including *all* of its different behaviours. Thus, they are still concerned with the very complexity from which we seek to free them.

The approach that appears to minimise disharmony most effectively is that taken in the design and provision of *orthogonally persistent object systems*, and particularly *persistent programming languages*. In this approach, well tried language design principles are adhered to, but they are applied to deliver a *complete computational environment*.

The term *persistence* is used variously in common parlance, but has been defined in this context to mean supporting data values for their full life times however brief or long these may be. The *life times* of data values are the period from their creation until they are no longer used by the persistent application. This range is illustrated in Table 2.

Table 2: The Various Life-times of Data Values

1	Transient results in expression evaluation
2	Local variables
3	Global variables and heap items
4	Data that lasts a whole execution of a program
5	Data that lasts for several executions of several programs
6	Data that lasts for as long as a program is being used
7	Data that outlives a succession of versions of such a program
8	Data that outlives versions of the persistent support system

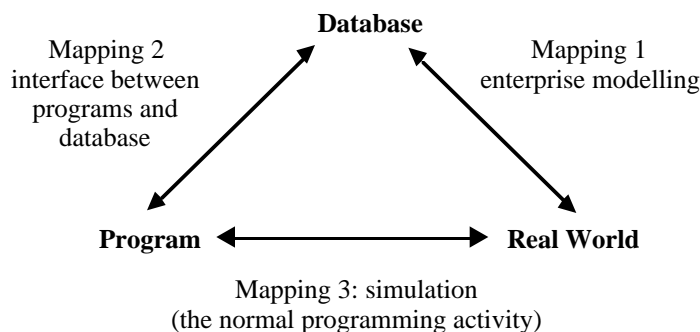
Typically, rows 1 to 4 have been serviced by programming languages and rows 5 to 8 by databases and file stores. Perversely, a barrier also exists between the forms of the data used. Short-term persistent data is often presented to the user as a collection of base types which may be aggregated or composed into constructed forms such as records, variants, functions and abstract data types. On the other hand, long-term persistent data often takes the forms of byte streams if held in a file system or some other structure such as relations in databases. Orthogonally persistent systems are designed so that the treatment of data values is uniform and independent of their longevity, size and type. Their goal is to achieve this uniformity for all aspects of the system services, from data definition and operations, to integrity, concurrency and distribution.

It should be noted that this goal for the provision of orthogonal persistence is independent of the choice of data model, type system, concurrency control mechanism, etc. that defines which values may occur and which operations may be applied to them. These choices are made for other reasons guided by their appropriateness for the application domains. However, there are some constraints on this choice if a safe system is to be provided for the PAS builder. Also, the engineering challenges in building such systems are such that these choices are not entirely independent of one another. This issue is investigated in later sections of the paper.

Similarly, the choice of programming paradigm (query language, imperative language, logic language, object-oriented language or functional language) is also orthogonal to the issue of providing persistence. Once again, appropriateness for the application will guide the choice.

The simplification achieved by orthogonal persistence is summarised in Figures 3 and 4.

Figure 3: Three Mappings with Two Support Systems



In Figure 3, a fragment of a PAS is represented built out of a combination of programming language and database facilities. Here the programmers have to build and understand two models of the same external (real world) system that the PAS

supports. Furthermore, they have to ensure that the view and behaviour of these models remains precisely consistent when accessed directly via database facilities and when accessed via programs using the database. This maintenance of precise consistency under all circumstances, including stress, is extremely hard for the application programmers to achieve. Either they achieve it with immense effort or they fail and the end-users suffer a less comprehensible system. There are other costs when construction is based on two subsystems. Programmers have to write translations between the two systems and computers have to execute them. Often this proceeds via the lowest common denominator of base types, losing structural information and protection.

Figure 4: One Mapping on an Integrated Support System



Mapping 1: simulation
(the normal programming activity)

In Figure 4, the same PAS fragment is envisaged constructed in an orthogonal persistent system. As the data is supported consistently whatever happens, the programmers only require to understand one model and maintain one mapping. The reduction in complexity for the programmers is obvious. It should also present a more comprehensible PAS behaviour to end-users.

2 Principles of Persistence

In order to recognise when good design has been accomplished it is necessary to have some design principles against which to judge the result. In constructing persistent systems the persistence research community has taken the view that the integrated approach pioneered by language designers should be extended to be more computationally complete. We regard the necessity of using some subsystem that is not specified (or specifiable) in the language as a failure of computational completeness. For example, using a file system is such a failure, since its semantics varies with the context of a program. These principles are developed by extending the design principles that work well for programming languages to encompass the requirements of persistence.

2.1 Design Principles

McCarthy [McCarthy *et al.*, 1962], van Wijngaarden [van Wijngaarden *et al.*, 1969], Strachey [Strachey, 1967] and Tennent [Tennent, 1977] all observed that expressive power in programming languages could be gained by separating the concepts and allowing them to be combined by powerful composition rules. Strachey and later Tennent distilled these ideas into three principles for use in the design of programming languages:

- the principle of correspondence,
- the principle of abstraction and
- the principle of data type completeness.

The principle of correspondence states that the rules governing the use of names and bindings in a programming language should be consistent. In particular the rules for introducing names and bindings in declarations should have a corresponding mechanism for abstraction parameters. This ensures that formal parameters behave consistently with local declarations. The principle of abstraction states that for all significant syntactic categories in the language there should be an abstraction mechanism. This allows essential details to be ignored by concentrating on the general structure. An abstraction consists of naming the syntactic category and allowing it to be parameterised. The most widely used forms of abstractions are functions and abstract data types. The principle of data type completeness states that

any combination or construction of data should be allowed in all types. As a consequence all data objects in a language should have the same “civil rights”.

The overall goal of the above principles is to design languages that are both simple and powerful. They are simple in that there are a minimum of defining rules with no exceptions, since for every exception to a rule the language becomes more complicated in terms of understandability and implementation. The minimisation of defining rules without exceptions also contributes to the power of the language since every exception makes the language less powerful in that it introduces a restriction. The expressive power therefore comes from ensuring that the composition rules are complete and minimal with no exceptions.

In persistent systems there are three design rules which have their origins in the above work. They are parsimony of concepts, orthogonality of concepts and completeness of the computational environment. The combination of these three rules yields the integrated persistent systems alluded to earlier.

Constructs are required that match well the application domain. The required parsimony of concepts focuses attention on to those essential constructs. The absence of extra features ensures simplicity. Orthogonality ensures that the system is composed of atoms, including persistent data, that may be combined in powerful ways to yield the appropriate abstractions. The case for completeness is that if the language is not computationally complete then the implementors and designers will need to combine the language with other facilities to build a PAS. This re-introduces the problems of disharmony.

Constructing persistent systems is made considerably easier when the whole computational environment is persistent. In such an environment programs and processes may be regarded as data and manipulated in the same manner, allowing transformations traditionally regarded as being performed by a separate mechanism to be executed within the persistent environment. For example, in the case of procedures, the persistent store may be used to provide a uniform library structure in a manner similar to providing a library of data parts [Atkinson & Morrison, 1985, Atkinson *et al.*, 1993, Kirby *et al.*, 1994].

2.2 Persistence Principles

The general principles presented above lead to specific principles concerning persistence.

2.2.1 Definition of Persistence

The persistence of a data object is the period of time for which the object exists and is usable [Atkinson, *et al.*, 1983]. We aspire to systems where the use of data is orthogonal to its persistence.

2.2.2 Orthogonal Persistence

There are three principles of persistence that may be used to achieve the above design goals. They are:

The Principle of Persistence Independence

The form of a program is independent of the longevity of the data that it manipulates. Programs look the same whether they manipulate short-term or long-term data.

The Principle of Data Type Orthogonality

All data objects should be allowed the full range of persistence irrespective of their type. There are no special cases where objects are not allowed to be long-lived or are not allowed to be transient.

The Principle of Persistence Identification

The choice of how to identify and provide persistent objects is orthogonal to the universe of discourse of the system. The mechanism for identifying persistent objects is not related to the type system.

The application of these three principles yields ***Orthogonal Persistence***.

2.2.3 Persistence Independence

Persistence independence frees the programmer from the burden of having to explicitly program the movement of data among the hierarchy of storage devices and from coding translations between long-term and short-term representations. As an example of persistence independence consider a sorting procedure that takes as input an array of objects to be sorted. The parametric array may be small or large and held in main store or disk. Beyond having to identify the correct array the programmer need not be concerned with the size or storage details of the arrays. Thus the user does not have to, indeed cannot, program to control the movement of data between long-term and short-term store; this is performed automatically by the system. The mechanical cost of performing the movement of data does not disappear but the intellectual cost does. That is, the programmer need not specifically write code for it, making the application code smaller and more intellectually manageable. The implementor of the support system now has the challenge of automating that data movement and any translation efficiently.

2.2.4 Data Type Orthogonality

Data type orthogonality is an aid to data modelling in that it ensures that the data model can be complete and independent of the persistence of the data being modelled. For example, bulk data types abstract over size and are therefore commonly used in persistent programming languages to aid the manipulation of massive collections of data such as scanned data from satellites or insurance policies sold by a company. Where such data is only considered long-term then the data model has to allow explicit conversion between long and short-term forms to allow creation of new bulk data and the manipulation of extracts from the long-term bulk data as short-term data.

Data type orthogonality also includes the language design principle of type completeness. For example, for bulk data this means that the elements of the bulk constructor are independent of the persistence of the data. Thus the programmer is not faced with a system where a set of one particular element type is allowed to be persistent whereas another element type is not.

2.2.5 Identification of Persistent Objects

A number of methods have been investigated to identify persistence of data. Some involve associating persistence with the storage allocator, variable name or the type in the declaration. Under the rule of persistence independence these are disallowed. They are also inappropriate for other reasons [Atkinson *et al.*, 1986].

In languages where the extent of a data object can differ from its scope, a limited form of persistence already exists. Such objects are normally kept on a heap and whether they can still be used depends on the availability of legal names. That is, they are kept available as long as the program and execution environment contains enough information to refer to them. This limited persistence can be extended to allow objects to persist beyond the activation of the program.

The technique which extends the above and is now widely used to implement this principle is *identification by reachability*. In this, the identification of persistent objects is performed by the system automatically by computing the transitive closure of all objects reachable (by following pointers) from some persistent root or roots [Atkinson *et al.*, 1982, Atkinson *et al.*, 1983, Cockshott *et al.*, 1984, Brown & Cockshott, 1985, Brown & Morrison, 1992]. Such data can then persist over

activations of the programs that operate on them. The analogy with automatic garbage collection is obvious.

2.2.6 Loss of Orthogonality

Loss of orthogonality of persistence occurs by disregarding any of the three principles given above. Most serious is persistence independence, as if it is broken it is hard to see how the persistence in the system is orthogonal in any way.

Some programming language designers have taken a very pragmatic approach to persistence and disregarded data type orthogonality for specific types. Pascal/R [Schmidt, 1977], DBPL [Schmidt & Matthes, 1992] and RAPP [Hughes & Connolly, 1990] are persistent programming languages based on Pascal where only first-order relations that contain no pointer types may be persistent. This corresponds to relational database practice but causes difficulty when other modelling techniques are used. The restrictions imposed are not necessary as it has been shown elsewhere that persistence can be added orthogonally to Pascal [Berman, 1991].

Where reachability is not used, languages often associate persistence with type. This immediately violates the principle of persistence identification and as a side effect, the other two principles as well. It also gives rise to dangling reference problems, or at least invalidated references, for persistent objects that point to non-persistent objects. The E programming language [Richardson & Carey, 1989, Richardson & Carey, 1990] most persistent extensions to C++ and the PGraphite language [Wileden *et al.*, 1988] all use this technique. Figures 5 and 6 illustrate this problem.

Figure 5: Persistent Objects before being sent to the Persistent Store

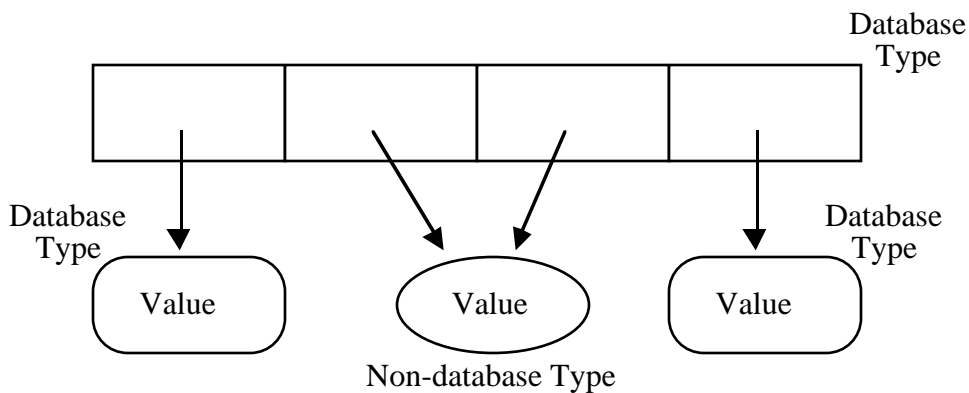


Figure 5 shows an object before being preserved in the persistent store. It has a type which permits its residence in the database and it holds two references to other objects whose types are such that they may also be stored in the database and another value which has a type for which database residence is not supported.

Figure 6: The Same Objects as in Figure 5 in the Persistent Store

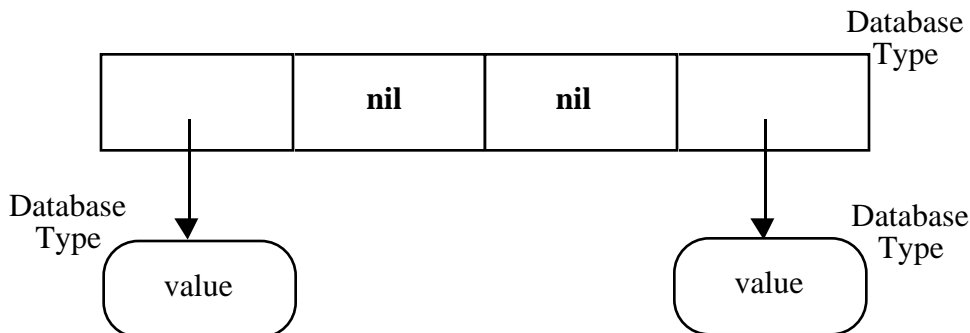


Figure 6 shows the result of saving Figure 5's structure in the persistent store. Only some of the structure is preserved and the remaining, non-database value is lost, being replaced by a dangling references or **nil** values if they are available.

Here the persistent store model of data is not consistent with the main store model, a position we wish to avoid since this adds complexity to the system for the programmer to master. For example, to overcome this loss of data when data is required to exist between program executions, programmers would need to do the following:

- 1 define a surrogate structure to the one shown that only used types that are supported by the database;
- 2 write code to translate the non-database values to equivalent database-values in this surrogate; and
- 3 write code to perform the inverse translation.

They would then have to ensure that this code was run at appropriate times and that the combined translation was exactly consistent.

The lack of completeness in the persistent computational environment can be seen in some C++ OODBMS. In these, the methods are held in a traditional program library and linked into C++ programs using a standard linker. The data values of the objects are held in the database. Garbage collection of unwanted data now becomes a problem in this model, since it is not known whether methods contain references to objects or whether any objects still exist that require the methods.

Some systems do not provide garbage collection of the data at all: a solution which will eventually lead to grief. Other systems allow explicit deletion of an object or recursive deletion from a root. In either case methods may still exist in the program library which now do not have any data; this usually causes run-time failures.

2.3 Savings with Persistence

The benefits of orthogonal persistence have been described extensively in the literature [Atkinson, 1978, Atkinson, et al., 1982, Atkinson *et al.*, 1984, Atkinson *et al.*, 1985, Atkinson and Morrison, 1985, Morrison *et al.*, 1985, Morrison *et al.*, 1986, Atkinson & Buneman, 1987, Morrison *et al.*, 1987, Atkinson *et al.*, 1988b, Dearle & Brown, 1988, Brown, 1989, Connor *et al.*, 1990, Cooper, 1990a, Cooper, 1990b, Albano *et al.*, 1993, Connor *et al.*, 1993, Morrison *et al.*, 1993, Morrison *et al.*, 1995]. They can be summarised as:

- improving programming productivity from simpler semantics;
- avoiding *ad hoc* arrangements for data translation and long-term data storage;
- providing protection mechanisms over the whole environment;
- supporting incremental evolution; and
- automatically preserving referential integrity over the entire computational environment for the whole life-time of a PAS.

The first saving of persistent systems is in the reduced complexity for application builders. Traditionally, the programmer has to maintain three mappings among the database model, the programming language model and the real world model of the application, as was shown in Figure 3.

The intellectual effort in maintaining the mappings and overcoming the complexity of the support system distracts programmers from mastering the inherent complexity of the application. In a persistent system the number of mappings is reduced from three to one (see Figure 4), thereby simplifying considerably the tasks undertaken during PAS maintenance and construction.

Corresponding to the intellectual savings of persistence, there is also a saving in the amount of code required to maintain the mappings. This has been estimated to be at

least 30% of the total code for a typical database application [King, 1978]. The code that is unnecessary is concerned with the explicit movement of data between main and backing store, and that required to change the representation of the data for preservation and restoration. An example of the former is input and output code, and of the latter is code to flatten and reconstruct a graph before output and after input. Thus the size of the application code is reduced thereby producing savings throughout the software life cycle of the application.

The third benefit of persistent systems is that a single model of protection may operate over the whole environment [Morrison *et al.*, 1990]. In most programming languages the simplest way to break the protection system is to output a value as one type and input it again as another. Thus security is lost over the persistent store. Using a single enforceable model reduces complexity while increasing the protection as its type-checking prevents disallowed operations and hence misuse.

Database programmers often build large systems by incrementally designing sub-systems and then by adding data and schema to the existing system. Initially a schema is defined, then the database is populated with data, and finally the programs (queries) are added. This is a powerful software engineering technique — that of incremental design and implementation. In persistent systems the paradigm is preserved but with rather more flexibility in the roles of the schema, data and programs. In particular, the schema is no longer considered to be relatively static and small, with the data and programs being dynamic and large. In persistent systems the schema, data and programs are all considered equally and may be static or dynamic, large or small, depending on the PAS under construction.

The final saving with persistence is that referential integrity of objects is automatically enforced. The referential integrity of an object means that, once a reference to an object in the persistent environment has been established, the object will remain accessible via that reference for as long as the reference exists. Furthermore, the identities of the objects are unique, and comparison of identity yields the same result independently of when it is performed. In a strongly typed persistent environment this also means that the type correctness of all such references is maintained, that is once a reference has been established, the type of the object referenced will not change. As will be seen later (section 5.2) this has a number of consequences for using references in source descriptions in PASs.

The provision of orthogonal persistence has several consequences for support systems. For example, much that is done explicitly by programmers in a non-persistent system has to be automated by the implementors of persistent support systems. The abstraction over longevity also requires the accommodation of schemata, programs and data of any size, since these may grow unpredictably during the lifetime of a PAS. An important aspect of such systems is the mechanisms which are used to construct sub-systems. Because of the reliability of control that is guaranteed by the persistent environment, such control can be delegated to tools without loss of safety [Connor *et al.*, 1994b].

Considerable research has now been devoted to the concept of persistence and its application in the integration of database systems and programming languages [Atkinson, 1978, Atkinson, *et al.*, 1983]. A number of persistent systems have been developed including: PS-algol [Atkinson, *et al.*, 1983, PS-algol, 1988], Abstract Data Store [Powell, 1985], Flex [Currie, 1985, Stanley, 1986, Stanley & Drummond, 1988], Galileo [Albano *et al.*, 1985], Amber [Cardelli, 1986], Trellis/Owl [Schaffert *et al.*, 1985], TI Persistent Memory System [Thatte, 1986], Tycoon [Matthes & Schmidt, 1992], Napier88 [Morrison *et al.*, 1994] and Fibonacci [Albano *et al.*, 1995]. The underlying goal in this work is to achieve engineering that is of sufficiently high quality to ensure that the orthogonal persistence abstraction performs well irrespective of the longevity and type of data and programs so that application programmers can program without having to think about data transfers and similar issues.

3 Integration Concepts

The integration of programming languages and databases into a single computational entity involves a number of challenges for researchers in identifying and unifying the concepts used for similar purposes in each domain. Of particular interest are: the unification of type systems and data models, binding mechanisms, techniques for identifying persistent data, system architectures and concurrency control, and taking advantage of the controlled environment. Progress in these areas is discussed below.

3.1 Types and Data Models

The long-term goal of research into persistent type systems is to unify type systems and data models by developing an adequate model of type that meets the computational needs of persistent systems [Connor, 1990]. Ideally we would like a simple set of types, and a type algebra, so that by a succession of operations and the provision of parameters, any data model or conceptual data model can be defined. This we have called the *type alchemist's dream* [Atkinson & Morrison, 1986].

Type systems provide both data modelling and protection facilities within databases and programming languages. Data modelling is performed in databases using data models, which have types to describe the form of the data, and in programming languages by using a classical type system. In both cases the universe of discourse of the system is defined by the set of allowable types which in turn are denoted by the set of legal expressions in the language. Data protection is provided by enforcing explicit and implicit integrity constraints in databases and by type checking in programming languages.

As a first step in the unification of data models and type systems some *approximate* equivalences can be recognised. These are summarised in Table 3.

Table 3: Equivalences between Data Models and Type Systems

Databases	Programming Languages
data models	type systems
schema	type expression
database	variable
database extent	value

While the equivalences are only approximate they do provide some insight into why integration may be possible at a conceptual level.

The issue of type checking is central to a type system that will provide data modelling and protection for persistent systems. Generally, data models in databases are concerned with the manipulation of the data that is consistent with the constraints imposed by the data model. In some cases these constraints may depend upon values calculated during the computation. As such they can be dynamic in nature and require dynamic integrity constraint checking for enforcement. By contrast classical type systems for programming languages are concerned with static checking which allows assertions to be made and even proved about a computation before it is executed. Static checking therefore provides a level of safety within the system. It also allows more efficient code since type checking code is not required at run-time.

At first the dichotomy between the checking times in databases and programming languages appears to be beyond resolution. One way forward is to pursue the limits of static checking while still accommodating the dynamic checking required in particular instances. This retains the safety and efficiency of static typing provided in programming languages in most cases while accommodating the dynamic flexibility of data modelling.

There are two approaches currently used to provide static checking for persistent type systems. The first is constraint specification where constraints over the data for a particular computation are expressed in some language. The checking requires a powerful theorem prover sometimes beyond the limits of those currently available. Such systems are usually undecidable and an unsuccessful check may be caused by the limitations of the theorem prover rather than inconsistent constraints. However, where the theorem prover fails it may sometimes provide useful information that can be used to form minimal dynamic checks [Sheard & Stemple, 1989]. This clearly delimits the points of dynamic checking while retaining as much static checking as possible.

The second approach is to extend classical type systems with specific types. As will be demonstrated later, points of dynamic checking are required to bind to persistent objects where programs and data are prepared separately and combined dynamically. The goal is to provide for some dynamic type checking while retaining as much static checking as possible.

Both of these approaches need to be pursued but only the latter is presented here in the next section.

One final difference is the semantics associated with the time of checking. In programming languages, the type checks, both static and dynamic, precede operations whereas dynamic database integrity constraint checks may be performed after operations but before commitment. In all cases, however, the strength of the checking is not compromised by the time at which it is performed.

3.1.1 Persistent Type Systems

The major challenges for type systems for persistent programming are to provide the richness of expression of data models and the completeness of protection required, in a mostly static form. These two goals are pursued in this section.

3.1.1.1 Programmer-defined Type Constructors

An important innovation in type systems has been the provision of programmer defined type constructors [Milner, 1978, Cardelli & Wegner, 1985, Cardelli, 1989]. This allows programmers to identify and introduce new and succinct notations corresponding to frequently used structures. Perhaps, more importantly, it also gives those structures a name that signifies their meaning to other programmers thereby fulfilling one requirement of types or DDLs: to describe data effectively.

A sequence of examples of declarations is given to convince readers that programmer defined types provide a convenient notation and an *extensible* descriptive system of comparable power to many data models.

The examples are written in a style close to the Napier88 notation [Morrison *et al.*, 1989, Morrison *et al.*, 1990] with the assumption that structural equivalence checking is in operation [Connor *et al.*, 1990]. A new type name is introduced with a notation such as the following:

```
type Count is int
```

which declares *Count* as a type name and makes it a synonym for the type **int** assumed already defined. Similarly, the following might define a *Date* type.

```
type Date is record (day, month, year : int)
```

These definitions become more interesting when type parameters are introduced:

```
type Pair [T] is record (first, last : T)
type DateStampedValue [T] is record (value : T; date : Date)
```

The first definition, *Pair*[T], describes an infinite class of record types with two labels, *first* and *last*, both of the same type. The definition of *Pair* may be parameterised by a type to yield a concrete type such as *Pair*[**int**].

Strictly these latest definitions are not types but type operators that when parameterised by a type yield a type constructor. We will make little of this distinction at the moment.

The second type definition illustrates how a type defined earlier may be used in another definition to build up a succession of more complex (more semantically meaningful) names of structures. In this case, *DateStampedValue* will generate record types that combine a date stamp and a value, for example:

```
type DateStampedImage is DateStampedValue [image]
type DateStampedCount is DateStampedValue [Count]
```

declares the type *DateStampedImage* whose instances would hold images and the date on which they were captured whilst instances of the type *DateStampedCount* would hold *Counts* and the date on which the count was made.

As a further illustration, the next two statements define co-ordinate types in integer and real spaces.

```
type IntXY is Pair [int]
type RealXY is Pair [real]
```

Regular structures such as those that occur in bulk types can also be defined. For example:

```
type Sequence [Element] is ...
type Ring [Element] is ...
type Set [Element] is ...
type Bag [Element] is ...
type Map [Domain, Range] is ...
type Tree [Key, Value] is ...
```

Readers have been spared the details of the definitions, typically they would involve abstract data types as described in the next section. These definitions can themselves be composed.

```
type MultivaluedFunction [Dom, Range] is Map [Dom, Set [Range]]
type Forest [Key, Value] is Sequence [Tree [Key, Value]]
```

They can then be used with more conventional definitions.

```
type Student is record (matric: int; fname, sname: string; dateOfBirth: Date; ...)
type Course is record (title : string; prerequisites : Set [Course])
type Enrolment is record (student : Student; course : Course)
type Curriculum is record (student : Student; completed : Sequence [Course])
type Class is Set [Student]
```

All the type operator names defined above: *Pair*, *DateStampedValue*, *Set*, *Bag*, ..., *MultivaluedFunction* and *Forest* can be used just like those supplied by existing data

models or those supplied initially in the language. Similarly, the types and constructors that have been produced and named, viz: *Count*, *Date*, *DateStampedImage*, *DateStampedCount* and *Class* can be used to produce instances just like types in a traditional data model.

It is therefore easy to envisage persistent support systems being shipped with few built-in types and constructors, but with extensive libraries of additional types and corresponding manipulation functions. Like traditional libraries of procedures they are then available for those who want them. This gives a simple starting point for application builders, but a potentially rich repertoire of additions which can be augmented further when needs have been identified, either by persistent support system suppliers or by the application builders themselves.

Use of all of the parametric type constructors described in this section is made much easier if associated libraries of generic procedures are provided [Atkinson, et al., 1993]. This is only possible if the language provides appropriate forms of polymorphism such as those described below [Morrison *et al.*, 1991].

3.1.1.2 Information Hiding using Existentially Quantified Types

The parameterised type operators introduced in the previous section define an infinite class of types. However, values can only be created with a specific type and the information contained in them is only available to those parts of programs that have the correct type descriptions, i.e. those type descriptions which are structurally equivalent. Two values created using a type operator are only type equivalent if they are created using equivalent parametric types.

The type operator mechanism can also be used to provide information hiding or abstract data types. By encapsulating the structure of the data behind a published interface, the type of the data, and thus its internal structure, may be hidden. In consequence, *two abstract data types with equivalent interfaces are equivalent irrespective of their internal type (or structure)*. This is best illustrated by an example. We start from a record type defining counters as a value and an increment operation.

```

type Counter [t] is record (value : t ; inc : proc (t → t))

let intInc = proc (x : int → int) ; x + 1

let intCounter = Counter [int] (0, incInt)           !create a Counter
...
intCounter.value := intCounter.inc (intCounter.value) !increment it

```

This creates a record with two fields, an integer *value*, and a function *inc* to increment the value. The type of the record is equivalent to

```

record (value : int ; inc : proc (int → int))

```

While the record is useful for aggregation it does not protect the data from outside interference. For example, both of the fields may be accessed and perhaps updated by any other part of the system that can specify an equivalent type. The final statement in the above example shows the *value* being updated by using the fields consistently as intended. However the type *Counter* does not restrict the use of the *value* to the function *inc*.

Second-order information hiding [Cardelli and Wegner, 1985] may be used to abstract over the type of the data and thereby the required consistent usage can be statically enforced. An abstract data type may be formed from a parametric type operator in the same manner as is used to create a concrete value. In this case, however, the parametric type (or witness type as it called in this context) once

instanciated becomes abstract. Consider for example the following revised creation of a counter.

```
let absIntCounter = abstract Counter [int] (0, intInc)
```

This creates an abstract data type, *absIntCounter*, where the two fields both operate over the same type but once the abstract data type is created, that type (the *witness* type) may never be re-discovered. As a consequence, only the functions within the abstract data type may operate on the value, for they are the only ones with the witness type. Since the witness type is abstracted over, but is known to exist, the type of *absIntCounter* can be written more formally as

$$\exists t. \mathbf{record} \text{ (value : } t \text{ ; inc : } \mathbf{proc} \text{ (} t \rightarrow t \text{))}$$

A second advantage is that all abstract data types formed in the same manner are type equivalent. For example

```
let absRealCounter = abstract Counter [real] (0, incReal)
```

also has the type

$$\exists t. \mathbf{record} \text{ (value : } t \text{ ; inc : } \mathbf{proc} \text{ (} t \rightarrow t \text{))}$$

since the witness type is abstracted over. These types may be used as follows, where *absCounter* is the type of all *Counters* as written in the above two examples.

```
let updateValue = proc (aCounter : absCounter)  
    aCounter.value := aCounter.inc (aCounter.value)  
...  
updateValue (absIntCounter)  
updateValue (absRealCounter)
```

One final problem remains for abstract data types. In order to preserve static type checking it is necessary to create a context in order to use the components of the abstract data type. The necessary restriction is that it must be possible to identify by static analysis that when the interface components are used, then they are only used in a manner that ensures that they all belong to the same abstract witness. Once this has been achieved with an appropriate syntactic mechanism, the requirement that the hidden information be only used consistently can be enforced by static analysis of the program text during type checking. This provides safety as it allows programmers to place precise limits on the ways in which long-lived data may be used. The details of models of witness types can be found in [Cardelli & MacQueen, 1988].

3.1.1.3 Generic Re-usable Persistent Polymorphic Procedures

The introduction of a statically typed universe trades the power of expression for static safety. How much the system designer is willing to trade depends upon the application domain.

In [Albano *et al.*, 1989] and its companion papers, an analysis of what constitutes a persistent type system is given. For modelling purposes it is generally agreed that some form of polymorphism is required to capture the expressiveness of data models and to increase component re-use [Morrison, *et al.*, 1987] as has already been demonstrated by abstract data types. The most favoured forms of polymorphism are universal polymorphism: parametric or inclusion.

Parametric polymorphism describes the polymorphism found in ML [Milner, 1978] and its derivatives whereas inclusion polymorphism is the style of polymorphism found in object-oriented languages such as Simula67 [Dahl & Nygaard, 1966]. An

interesting hybrid may be found in the database programming language Galileo [Albano, et al., 1985], which is a derivative of ML but utilises inclusion polymorphism to implement part of the Semantic Data Model [Hammer & McLeod, 1981]. Cardelli [Cardelli and Wegner, 1985] has shown separately how the parametric and inclusion forms of polymorphism may be integrated, as bounded quantification, to yield forms of abstraction not available to either one.

Since parametric and inclusion polymorphism give similar power [Cardelli and Wegner, 1985] only parametric polymorphism is described here. In parametric polymorphism, functions may be defined that work for all types. Consider a function that counts the number of elements in a list. Since it does not use the values in the elements of the list itself then the function should work for all lists. It is said to be universally quantified and its type may be written as:

$$\forall t.(\text{List } [t] \rightarrow \text{int})$$

To extend this example a little, consider functions that maintain an index. The index would store values of any type and use keys of any type which has (at least) equality defined over it. The functions therefore require to be polymorphic over the value and key types. This is relatively easily specified using parametric polymorphism. For example, if the index is implemented by a list of pairs, one element for the key and one for the value, then the type of the function to enter a value in the index is

$$\forall \text{Key}. \forall \text{Value}. (\text{Key} \times \text{Value} \times \text{List } [\text{Key} \times \text{Value}] \rightarrow \text{List } [\text{Key} \times \text{Value}])$$

The advantage of this type of abstraction is that the functions can be placed in the persistent store and used later to operate over indexes of all types. The software engineering advantages of writing concise code that is generally applicable are well known.

Further abstraction can be obtained by observing that the functions should be able to abstract over the structure of the data as well as its base type. The above function can be written since the polymorphic procedure heading specifies that the index is implemented by a list. Thus it may be implemented using the operations on a list. It is possible to implement the index using other structures such as arrays or B-trees. For example the type using a B-tree is:

$$\forall \text{Key}. \forall \text{Value}. (\text{Key} \times \text{Value} \times \text{Btree } [\text{Key} \times \text{Value}] \rightarrow \text{Btree } [\text{Key} \times \text{Value}])$$

Combining the functions requires that the parametric types are equivalent. Abstract types may also be used to yield a combination of universal and existential quantification. Note that with universal quantification one abstract polymorphic form can be written and special cases generated, whereas with existential quantification existing values are described by a more general type allowing more general abstraction over that type.

Polymorphism is one technique for retrieving some of the expressive power that is lost by the introduction of the type system in the first place. By abstracting over the types more general computations can be expressed than in monomorphic systems. Polymorphism does not regain all of the power lost by the introduction of a statically checked type system. Relational systems are highly polymorphic but depend upon the dynamic evaluation of types for their expression. For example, the specification and implementation of a generic natural join function provides an example of abstraction over types that is beyond the capabilities of most statically typed polymorphic type systems. This is because the details of the input types, particularly the names of the tuple components, significantly affect the algorithm and the output type.

Type systems that can accommodate functions such as natural join are generally higher order and dynamically checked. An interesting exception in this case is the Machiavelli type system [Ohuri *et al.*, 1989]. In general, however, in order to retain static checking and still write these higher-order functions a different technique is

required. One such technique is linguistic reflection [Stemple *et al.*, 1992b] discussed later in this paper.

3.1.1.4 Persistent Type Checking

We first demonstrate the need for some form of dynamic type checking in persistent systems and then show how it is provided.

The novelty of a persistent type system is that it must provide type checking over the whole of the computational activity including the use of persistent data. Within a traditional file system, the use of persistent data is achieved by programs having textual descriptions of how to find data within the file system. For type checking, a static description of the file data within the program is used by the compiler to assert that the data will be of the specified type when it is provided from the file. For implementation, the run-time system must perform dynamic type checks to enforce these assertions. Figure 7 illustrates the point.

Figure 7: Dynamic File Access

```
type Address is record (name : string ; age : int ; gender : bool)
var fileDescriptor : file [Address]

fileDescriptor := open ("/user/ron/addresses")
! At this point the system accesses the file ensuring that it exists and
! checks that the type of the data in the file is of the specified type

if fileDescriptor = nil then error ... else
begin
  while ~eof (fileDescriptor) do
    begin
      let this = read (fileDescriptor)
      ...
    end
  end
end
```

In Figure 7, a type *Address* is declared to be a record with three fields *name*, *age* and *gender* of types *string*, *integer* and *boolean* respectively. A variable *fileDescriptor* is declared as having the type of a file of *Address*. It is initialised to the file at path name *"/user/ron/addresses"* by the *open* statement. If the file exists, then a dynamic type check is performed to ensure that the data contained in the file is of the correct type, namely *Address*. Note that the interpretation of the path name string during the *open* operation is an example of the failure of computational completeness described above as it is not determined by the semantics of the language.

It should be noticed that in Figure 7 there is only one point of dynamic typing, that is during the *open* statement. The program is compiled with the assertion that the file will have the correct type and all other type checking can therefore be performed statically. Thus the program may read values from the file without performing subsequent dynamic type checks. In this form of persistence each file can be considered as a persistent root. Notice also that the type of the values in the file must be guaranteed by the file system. Thus, overwriting the data with values of a different type is forbidden for a strongly-typed file system.

The nature of the *open* statement in Figure 7 is worthy of closer inspection. The statement has different types depending upon the values contained within the file. In the above case it returns the type *file of Address* but in another case may return, for example, the type *file of Employee*. The most general type for *open* is one that can be coerced into the specific type dynamically, during the *open* statement. Thus *open* can be said to have a dynamic type.

Changing the type of the values in a file entails deleting the file and replacing it with a file of a different type. This appears to the program as if the old values have been coerced to a new type. Notice however that the coercions are all performed as a side effect of another operation and that there need not be an explicit dynamic type nor coercion facilities within the language. In a typical implementation, for example an Ada APSE, these operations depend on two systems with independent semantics which can be changed independently. A *deus ex machina* may therefore change the type or existence of the file under the feet of the programmer.

As the number of typed files in the persistent space grows, then programs will inevitably wish to access more than one file and to have inter-file references. This is the major difference between a typed file store and a typed object store in that the files are self contained with no inter-file references, whereas objects are highly inter-connected.

Let us assume for a moment that typed files could be extended to accommodate inter-file references. As a consequence of inter-connection from a file, that file's type must have knowledge of the types of all the files to which it connects. Over time the persistent space becomes highly connected necessitating every program to have a near complete description of the types of all the files.

The second problem with the typed file approach to persistence is that the file types are fixed at file creation time. Files cannot refer to newer files whose type has not yet been determined and this restricts evolution of the system. At a programming level, inheritance partially overcomes this by yielding additive evolution [Atkinson *et al.*, 1993], which allows values to be replaced by more specialised forms. Thus in Figure 7 the file could be one that provided an extra address field, like *postcode*. The processing system must be set up so that it ignores this to achieve correct execution. To accommodate evolution completely it is necessary to allow arbitrary changes in the data model, at least at some well defined points.

Both problems outlined above: partial specification of the type structure (schema) and the evolution of the data model can be solved by the same mechanism that performs the dynamic type checking of persistent data. The integrated solution involves an infinite union type of which type **pntr** of PS-algol [PS-algol, 1988], **dynamic** of Amber [Cardelli, 1986], **any** and **env** of Napier88 [Morrison, et al., 1994] are examples. Persistent roots have the infinite union type and values are injected into these objects with a type and projected dynamically onto a type for re-use. It is at the points of projection that the dynamic type checking occurs. Figure 8 illustrates this point using type **any** and the Napier88 notation.

Figure 8: Projection from an Infinite Union

```

type Address is record (name : string ; age : int ; gender : bool)
let ps = PS () !This is the only standard function in Napier88
                ! It returns a persistent root of type any
                ! Before the value can be used with its specific type
                ! it must be projected onto that type

project ps as X onto
Address : write X (age) ! X has type Address here
...
default :           ! This is a "catch all" and ps has type any here

```

In Figure 8, the standard function *PS* is called. It returns the value of the persistent store which has the dynamic type **any**. The use of a dynamic type allows the persistent store to change its specific type between activations of particular programs. In order to use the value with its most specific type, *ps* is projected by the **project**

clause onto that type. It is during the projection that the dynamic type check occurs. Within the projection context, the value may be used with its most specific type, in this case *Address*. The above program writes out the *age* field of the record. Any number of types may be specified in the projection; the first correct one is used. A catch all **default** clause is used to trap all the cases where none of the specified types match. Within this context the value only has the type **any**.

Thus the mechanism is that the persistent store has a most general dynamically checked type. To use the values within the store the dynamic type must be projected onto the specific type for the store. The method by which the store may change its type, by injection of a value with a different type, is not of interest here except for the fact that it can be performed.

This mechanism occurs implicitly in standard database interfaces. When a program opens a database it specifies a schema or view. During the opening operation the schema or view used during the program's compilation is compared with the database's current schema or view. Arbitrary changes may have been made to the database using a schema editor between the compilation and this execution. If the schema or view no longer matches the expectations established at compilation-time then an error is signalled. Thus, internally the run-time system is able to treat the database as having a dynamic type and to perform a dynamic verification that the expected and actual types match.

From Figure 8 it can be seen that dynamic checking is limited to the points of projection from the type **any**. Values of these dynamic types are first class and may also be a constituent part of any other type. Thus, a graph of objects in the object store may have one root of type **any** as well as other values of type **any** that are components of other objects. Before they can be used with their specific type they also have to be projected onto that type. This allows programs to specify only the part of the schema that they require up to a point of dynamic checking. As a consequence a schema, which is represented by an arbitrarily large collection of mutually referencing types, may scale well since the schema specification is bounded by the dynamic types. Incremental schema changes inject new values into an **any** and the type of the rest of the specification remains unchanged. In addition, where an **any** encapsulates the type, type checking is postponed until required. Hence excessive type checking costs on start-up are avoided. This is illustrated in Figure 9. To implement this delayed checking it is necessary to store, potentially large and complex types with the **any** value.

Figure 9: A Partial Specification of Types

```
type Address is record (name : string ; age : int ; gender : bool ; extra : any)
let ps = PS ()
project ps as X onto
Address :
  begin
    let this = X (extra)           ! this is of type any
    type extraInfo is record (idNo : int ; spouse : Address)
      ! Programs not using the extra field do not need to specify
      ! extraInfo
    project this as Y onto
      !type check using extraInfo only happens when this is
      ! executed
      extraInfo : write Y (idNo) !write out the id number
      ...
    default : ...
  end
default : ...
```

In Figure 9, where the *extra* field is used the specification of the *extraInfo* type is required. Where programs do not use the *extra* field the type *extraFieldInfo* does not have to be declared. Thus only part of the type structure need be specified, the part of interest to the program, as shown in Figure 10.

Figure 10: A Partial Use of Types

```
type Address is record (name : string ; age : int ; gender : bool ; extra : any)
let ps = PS ()
project ps as X onto
Address :
  begin
    write X (name)           !write out the name
  end
default : ...
```

The use of dynamic types allows the data model to evolve without recompiling all the programs that refer to the data. For example, if the *extraInfo* type is altered, then only programs that used that type need be altered.

To summarise, explicit dynamic types allow partial specification of the overall structure of the data (schema) and facilitate the evolution of the data, without having to alter programs that do not make use of the evolutionary changes.

3.1.1.5 The Range of Type Checking Times

The addition of the infinite union type **any** facilitates incremental dynamic checking in persistent systems. Thus the range of checking times includes just-before-use dynamic checking. At the other end of the spectrum the presence of the persistent store allows persistent values to be bound to programs during program construction. Where the programs themselves are persistent objects, this leads to the concept of hyper-programming [Kirby *et al.*, 1992] that will be discussed later. For the present it is sufficient to realise that persistent values are available to the type checker allowing types to be value based. This leads to persistent systems being able to support some *dependent types* (i.e. types which require tests on values to establish their

equivalence) similar to static constraints based on values in traditional database systems.

To introduce the usefulness of dependent types, a motivating example of dependent bulk types taken from [Connor, et al., 1993] is given. The semantics of many bulk type data models depends on user-defined attributes such as definitions of element equality, ordering, and other domain predicates. While these attributes are an intrinsic part of the data model, they are not normally treated as part of the type description. This may lead to the occurrence of data modelling errors such as a union operator accidentally being applied to two sets which have different semantics for the equality of their elements.

Directories indexed by Scottish names provide an example where the inclusion of element attributes in the type may be a requirement. People may take a different view as to whether the names “MacFarlane”, “Macfarlane”, and “McFarlane” are really different names, but their owners are usually protective of their different forms. If they are used to retrieve data however, it is unlikely that the retriever will wish to distinguish between them. There is therefore a requirement for different types of directory dependent upon the domain equality semantics. For two directories to have the same type they must have (depend on) the same equality operator, which is a value. If it is a persistent value then they can depend on the same value even if they were formed on different occasions.

For dependent types to be structurally equivalent it is necessary for the values upon which they depend to be equal. For the equivalence of dependent types to be statically decidable it is therefore necessary for the values upon which the types depend to be statically available to the type checker.

The essential requirement is that any values upon which a type depends are evaluated before any equivalence testing is performed upon the type. This restriction of the general type description can be enforced by restricting the dependencies to being existing persistent values.

A full description of such a language mechanism, including the explanation of the necessary restrictions and an introduction to the subject of polymorphism over dependent types, is given in [Connor, et al., 1993].

3.2 The Range of Binding Mechanisms

Binding mechanisms present the user with a trade-off between safety and flexibility [Morrison *et al.*, 1987, Morrison *et al.*, 1990]. Dynamic binding is the most flexible since the binding is delayed until the latest possible time at which a choice can be made. In contrast, static binding is safer, in that static checking may be used to give advice to a person better able to understand it (a programmer) and to eliminate run-time binding errors. The programmer has to choose the mechanism most suitable for a particular application from the range of binding available within the construction system. Persistent systems extend the range of binding by presenting the user with the possibility of using persistent objects during the construction of an application.

Traditionally, in programming languages and database systems, a binding occurs between a name and a value [Strachey, 1967]. That is, a value is bound to a name for some period during the evaluation of a program or query. This has been extended by Burstall & Lampson to include a type [Burstall, 1984] and further by Atkinson & Morrison to mutability [Atkinson & Morrison, 1987]. A binding mechanism, therefore, has four components: a name, a value, a type and an indication as to whether the value is mutable or not. To complicate the issue further, bindings may be performed statically by the compiler or dynamically by the run-time system or indeed, as is often the case, at intermediate stages.

3.2.1 The Nature of Binding Mechanisms

As indicated above, the binding mechanism has four components:

- Is the binding to a mutable (L-value) or immutable (R-value) value?
- When is the binding performed?
- What scoping is involved?
- When is type checking performed?

Bindings may be made to immutable values, that is, constant values that do not alter during the period of the binding, or to mutable values (locations) where the binding does not change although the value referred to by that location may change. These kinds of bindings are traditionally known in programming language parlance as R-value (for immutables) and L-value (for mutables) bindings [Strachey, 1967]. The manifest constants of BCPL [Richards & Whitby-Strevans, 1979] or Pascal [Wirth, 1971] are examples of R-value bindings. On the other hand, Pascal variables are examples of L-value bindings where the compiler may bind the name to a location but not to a particular value since that may vary at run-time.

A binding is always performed with reference to a particular environment. That is, a binding will be part of a particular environment and may use other bindings in that environment to establish its own. The scope of the binding determines where it may be used. There are two common forms of scoping in programming systems, static and dynamic. In static scoping, the scope of the bound name can be detected by static analysis of the program. Algol 60, and all derived languages such as Pascal, Ada, etc., utilise a static scoping rule that allows duplicate bindings to be detected statically. In dynamic scoping, the binding in scope is the one that was last defined in the dynamic evaluation of the program. Dynamic scoping can be seen in Lisp [McCarthy, et al., 1962], the binding of file names and in the segment binding mechanism of Multics [Organick, 1972] and all derived operating systems.

Type checking can be performed statically by a compiler or by the run-time system. Dynamic type checking occurs when the run-time system executes code to ensure that the data is of the correct type. This typically occurs even in so called statically checked languages when external data is brought into the computation e.g. in read statements and, as has been shown earlier, during projections out of infinite unions.

A categorisation of binding is given in Table 4.

Table 4: A Categorisation of Binding

	Static R-value	Static L-value	Dynamic R-value	Dynamic L-value
Static Typing Static Scoping	1	5	9	13
Static Typing Dynamic Scoping	2	6	10	14
Dynamic Typing Static Scoping	3	7	11	15
Dynamic Typing Dynamic Scoping	4	8	12	16

There are 16 different methods of binding based on the four binding choices given above. The most static form is a static R-value binding with static type checking and static scoping. The most dynamic form is a dynamic L-value with dynamic type checking and scoping. It is interesting to note that even within one particular language there is often more than one binding category. For example, in Pascal category 1 describes **const** values, category 13, variables, category 15, variant projections (cases in Pascal) and category 16, file names.

In determining the appropriate binding mechanisms for a particular system, the designer is faced with the problem of balancing safety against flexibility in a similar manner to type checking. The safety in the system is derived from being able to say (even prove) something about the program before it runs (i.e. statically) in order to improve confidence that it is correct.

An aspect of static checking that is often overlooked in programming systems is that the source code then acts as better program documentation. If a compiler can statically check a program then so can a programmer. Thus statically checked programs have better documentation properties and consequently better cost properties throughout the life cycle of the programs.

3.2.2 The Range of Persistent Binding Times

It may be argued that the whole spectrum of binding mechanisms is required in persistent systems to facilitate the needs of prospective programmers. From observations the uses of bindings are:

- the creation of new objects (and binding into the persistent store);
- the reuse of existing objects (program and data) by new objects;
- new combinations of existing and/or new objects; and
- incremental construction of objects.

For example, the programmer may wish to bind statically to a combination of objects in the persistent store in which case the objects are assembled into a new object and bound together. On the other hand, the programmer may wish to obtain the latest version of the object by delaying the binding until the object is about to be used.

The presence of a persistent environment in which programs are compiled and where programs may refer to persistent objects gives rise to three new perspectives in binding. They are:

- the extended scoping of bindings;
- the need to break bindings; and
- the separation of names and values in persistent bindings.

The presence of a persistent environment in which programs are composed, compiled, linked and run extends the scoping of bindings to persistent values. Where the source programs themselves are values, static R and L-value bindings may be made in the program source to persistent values. Thus the source programs may contain already evaluated values that have been placed in the persistent store. These are manifest to the program and as such belong to category 1 which up until now in not much used in programming systems. The advantage of this style of binding has already been shown in dependent types and will be shown later to lead to the concept of hyper-programming: a new technology only available in orthogonally persistent systems [Kirby, et al., 1992].

Persistent systems also highlight the need to break bindings and to recreate them later or in some other context. Consider the problem of releasing systems. In order to ship a system from one environment to another it is often sensible to ensure that the system fits in with the new environment. In Unix this often involves redefining shell variables or late binding the system calls to the new environment. For example, an application that uses the Unix *malloc* system call is both unlikely and unwise to take its local copy of the command with it when it is shipped to a new environment.

The environment of a persistent value includes the graph of values which can be reached from it and shipping values from one persistent system to another may involve copying the complete graph. This is not usually desirable for the semantics of the value, the utility of the value in its new environment nor from an efficiency point of view. In order to accommodate this rebinding, bindings must first be broken which is in itself a dangerous activity. One proposal for achieving such rebinding in a safe

manner is that of Octopuses given in [Farkas & Dearle, 1993]. This allows the bindings within a value to be broken, producing a wiring diagram of the broken bindings. This wiring diagram may then be rebound in the new environment when the value is moved. Since the wiring diagram is effectively an abstract data type it may be programmed to achieve a number of varied rebindings. The mechanism is a special case of linguistic reflection addressed specifically to this problem.

Flexible binding is only possible where there is a separation of the concepts of names and values. For flexibility a value may participate in many bindings and have many names (aliases). More importantly, in persistent systems, the name space is separate from the value space allowing the graph of persistent values that constitutes the persistent store to have any number of name spaces layered on it. Application specific naming schemes are supported by this arrangement.

It should be noted that polymorphism allows the separation of values and types, structural type equivalence allows the separation of names and types, and dynamic typing the separation of names and values. This is another example of the power of orthogonal design.

Some of the skill in using a persistent system will be in deciding which objects are statically composed and never changed and which objects are dynamically composed. A judicious mixture of mechanisms has to be provided by the system and eager binding is expected to be the preferred style of usage, in that, when it is appropriate to perform the binding, then it is done for safety and not delayed unnecessarily.

3.3 Concurrency

Traditionally the database and programming language communities have taken different approaches to concurrency control. In programming languages, concurrency control is based upon the concept of the co-ordination of a set of co-operating processes by synchronisation. Language constructs such as semaphores [Dijkstra, 1968b], monitors [Hoare, 1974], mutual exclusion [Dijkstra, 1968a], path expressions [Campbell & Haberman, 1974] and message passing [Brookes *et al.*, 1980] have been provided to support this concept. By contrast, in databases, concurrency is viewed as a system efficiency activity which allows parallel execution and parallel access to the data. However, each database process may have to suffer the indignity of abortion in order to sustain the illusion of non-interference. The key concept in databases is that of serialisability [Eswaran *et al.*, 1976] which has led to the notion of atomic transactions [Eswaran, et al., 1976, Kung & Robinson, 1982] supported by locking [Eswaran, et al., 1976] or optimistic concurrency control methods [Kung and Robinson, 1982].

In both cases the user must attempt to understand the computations in terms of some global cohesion. In programming languages the emphasis is on synchronisation and the overall cohesion is understood in terms of the conflation of all the synchronisations. A number of techniques, including CCS [Milner, 1980] and the π -calculus [Milner, 1991] have been developed to help with this. In database systems, global cohesion is understood in terms of the concept of serialisability [Eswaran, et al., 1976] but includes failure semantics such as abortion.

Figure 11, taken from [Munro *et al.*, 1994], illustrates a spectrum of understandability from the points of view of programming language and database users.

Figure 11: A Spectrum of Understandability

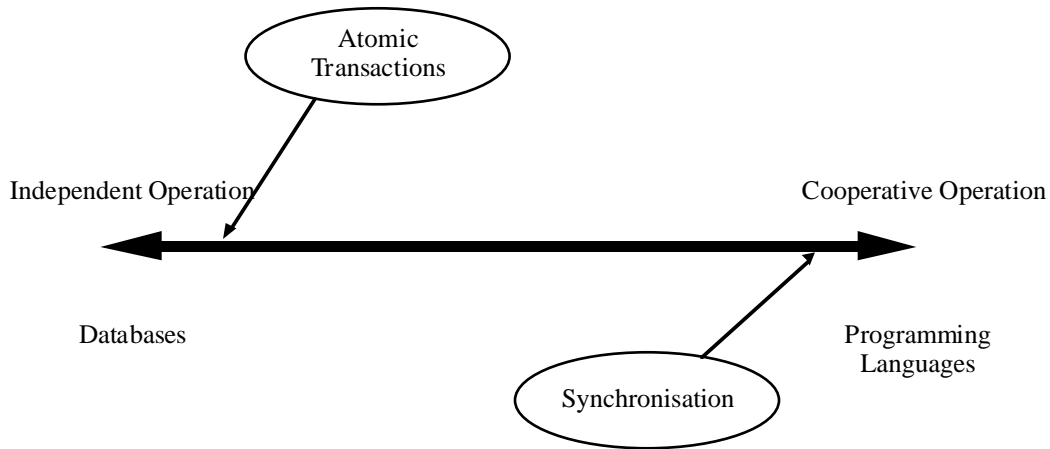


Figure 11 illustrates that databases tend to use atomic transactions to enforce isolation rather than co-ordinated sharing. Programming languages promote co-operation. Thus in integrating databases and programming languages, the designer must unify these established and provenly useful positions. The impetus does not altogether come from persistence however, since languages that support atomic transactions and databases that require non-serialisable and designer transactions [Ellis & Gibbs, 1990, Sutton, 1990, Nodine & Zdonik, 1992] have been identified as necessary by their respective communities.

Both camps agree that co-ordinating a set of computations that share data is a complex undertaking and that it is difficult to characterise the power and behaviour of the mechanisms and even more difficult to compare them with each other.

Previous attempts to provide concurrency in orthogonally persistent systems either have focused on a single model [Morrison *et al.*, 1988, Morrison *et al.*, 1989], provided a mechanism for extending synchronisation or provided basic persistent threads. The work of Krablin's CPS-algol [Krablin, 1987b, Krablin, 1987a] is perhaps the most notable. CPS-algol added constructs to PS-algol to support processes and hence persistent threads. The concurrency model is co-operative based on synchronisation by conditional critical regions. Concurrency is provided by executing procedures as separate threads. Using these primitives and the higher-order functions of PS-algol, a range of concurrency abstractions can be constructed including atomic and nested transactions as well as more co-operative models. This work has been re-visited recently in [Munro, 1993] and [Matthes & Schmidt, 1994].

Given that the overall goal is for the user to understand the computations in terms of some global cohesion, the CACS system [Stemple & Morrison, 1992, Morrison *et al.*, 1993] proposes a mechanism for integrating concurrency control in programming languages and databases. It takes the point of view that the difficulty stems from both the low-level nature of the mechanisms and the inherent complexity of the problem.

The goal of the CACS approach is to control the coherence of sequences of operations on shared data in an understandable and flexible manner. The essence of the system is: understandability; the separation of concurrency control from data; formal capture; and a path to implementation. The global cohesion of an action in CACS is visualised as the movement of data among access sets. The visibility of the data to other users then becomes an issue of synchronisation. Where the visibility coincides with commit time, atomic transactions may be obtained, and where the visibility is immediate then synchronisation is present. By viewing the co-ordination of the use of data as the co-ordination of the movement of the data among access sets, then the same data may be used in a different manner at different times. Thus the data may take part in an atomic transaction one day, a saga the next and may be in a co-operative computation the

next. A future goal of the work is to allow interaction of the concurrency control methods.

4 Technology to Support Persistence

An overview is now presented of some crucial aspects of the technology needed to support orthogonally persistence systems. The explanation below is not intended to be sufficient for someone trying to implement persistence though the citations are intended to give an entry to the pertinent literature. It is intended to make readers aware of what is involved and convince them that orthogonal persistence is feasible.

4.1 Implementation Architectures

To support persistent applications a stable and reliable store to hold the data and an execution mechanism to execute programs are required. The principal categories of supporting architecture are discriminated by the way in which these major components are provided and the way that they interwork. Four architectures are recognised in order of increasing commitment to the persistence philosophy. Not every reported system fits precisely into these categories and their state of development varies between commercial products and research prototypes.

4.1.1 Combined Existing Systems

The designers may start from an existing data model (e.g. relational, object-oriented, etc.) and combine it with an existing language (e.g. C, Fortran, C++, Cobol, etc.). The combination proceeds by arranging a notation to perform bindings. This involves identification of types and data via databases, schemata, views and queries and component use in programs or methods. The database component operations are made manifest by extending the language or providing some library.

The resulting combination may be sympathetic to the original language to varying degrees. Difficulties are often met because of the limited overlap between the two universes of discourse. Besides the type and conceptual differences to be overcome, syntactic noise is often introduced by attempting such a combination. There are, however, obvious advantages: capitalising on training, utilising well polished existing system code, continuing the use of legacy code or data with relatively little re-coding and permitting multi-lingual working against the database.

Where the language and data model are reasonably sympathetic, an acceptable quality of integration can be obtained, for example the integration of O_2 with C++ [Deux, 1990, Deux, 1991]. However, it would be unrealistic with this architecture to expect that all the features of the programming language would work directly on persistent data and that data of all types could be stored. The storage of large collections of objects within a standard relational system is an interesting example of this approach [Reinwald *et al.*, 1994]. There, the programmer can use the full facilities of C++ to build and operate on an arbitrary group of objects and use the full relational facilities to store and retrieve such groups. However, programmers are not permitted to store persistent references between groups.

Implementations taking this approach attempt to leave both the language run-time system and the DBMS unchanged. Typically they are run as separate processes that communicate via messages.

4.1.2 Extended Existing Systems

This approach takes place in one of two directions. Designers may start with an existing database and extend it to have more complete type and computational facilities, or they may start with a language and add persistence. Postgres is a good example of the former [Stonebraker & Kemnitz, 1991]. Commercial systems, such as Microsoft ACCESS [Microsoft Corporation, 1994a, Microsoft Corporation, 1994b] also fall in this category as they start with a relational model, and expand its computational facility by adding program and interface generation facilities. The

development of object-oriented facilities in the SQL3 standard [Kulkarni, 1994] corresponds to this approach.

The first version of PS-algol [Atkinson, et al., 1983] was an example of this second approach, as it began by extending S-algol [Morrison, 1979]. Similarly, Pascal-R was initially described as the addition of relational capabilities and long-term storage to Pascal [Schmidt, 1977]. There are, however, fundamental differences between these two examples: the former set out to leave the language's type system unchanged and hence to achieve orthogonal persistence, the latter set out to extend the type system, but made no attempt at data type complete persistence.

Many persistent systems continue to be developed using this second approach. Typically they are persistent versions of C or of C++ [Richardson and Carey, 1989, Shapiro *et al.*, 1989, Richardson and Carey, 1990, Reinwald, et al., 1994] but persistent versions of SmallTalk [Straw *et al.*, 1984], ML [Matthews, 1985, Matthews, 1989, Nettles & Wing, 1992] and Ada [Wileden, et al., 1988] have also been reported. Where the persistent language takes this form, there is the advantage of converting existing software to have persistent behaviour with minimal effort.

Construction of a persistent system by extension of an existing system usually requires significant modifications to the original system's run-time system. Discussion here focuses on the case where a programming language is extended. There are four tasks to automate:

- obtaining the type information from programs and accumulating it as a schema;
- storing, using a representation that will always be interpretable, any values that future programs may use, in a stable transactional store;
- identifying which values should be promoted to longevity and detecting when they can no longer be used so that their resources can be re-used;
- arranging to load into a program's active store any persistent data that it is about to be used.

Often the implementation involves new persistent store technology, including recovery, transactions and space management. There are variations regarding the extent to which the store technology is tuned to the requirements of the particular language supported [Brown *et al.*, 1992] and the extent to which the store builds-in clustering, indexes, bulk types, etc. [Cluet & Delobel, 1991, Matthes & Schmidt, 1991, Zezula & Rabitti, 1992, Cluet & Moerkotte, 1993].

Some persistent extensions to languages accept that they cannot implement all four of the tasks listed above. For example, the Texas store [Singhal *et al.*, 1992] provides a stable checkpointable memory in which C and C++ programs run. In consequence it provides orthogonal persistence for those languages as the programs run unchanged in this store, but the representation would not carry data across changes to parts of the support system and no automatic management of object lifetimes (task 3 above) is supported. These limitations are an ineluctable consequence of supporting any language that happens to be run in the store, as it means that there isn't a complete and reliable type system from which to derive the required information to provide these services.

4.1.3 Systems Based on an Integrated Design

Recently, several systems have been developed that are based on an integrated design, which tries to equitably use experience from both programming languages and databases, but starts afresh to develop a single system that provides the functionality of both systems. Examples are: Napier88 [Morrison, et al., 1994], Fibonacci [Albano, et al., 1995] and Tycoon [Matthes *et al.*, 1994] (all products of the ESPRIT Basic Research Action 6309: FIDE₂).

These generally have the advantage of being conceptually simple compared with the preceding approaches. However, they usually do not manage to provide full database facilities — that is, few can actually demonstrate a complete repertoire of incrementality, transactions, recovery, concurrency, distribution and scalability. (It *appears* that this is more a consequence of teams being unable to muster the effort to tackle all of these issues together rather than of fundamental limits.) These support systems have to provide an interface to legacy code as they are typically monolingual and have a unique data model or type system. Their implementation still has to achieve the four tasks identified above.

4.1.4 Persistent Worlds

All three of the categories discussed above are typically implemented above standard operating systems (UNIX, Windows NT, etc.). This has the advantage that it is easier to export the research to other sites. It has the disadvantage that it is extremely difficult to obtain entirely consistent behaviour and reasonable efficiency on these platforms.

This leads to an alternative line of research where a totally new computational platform is constructed. This may be a persistent operating system (on top of which, all languages achieve persistence automatically, and with which all data has consistent persistent behaviour) or even new hardware architectures with associated operating systems. Typical of the former are: Grasshopper [Dearle *et al.*, 1994], EOS [Gruber, 1992, Daynès & Gruber, 1994] (also a FIDE₂ product) and of the latter are: Rosenberg's proposals [Rosenberg, 1990] and DAIS [Russell *et al.*, 1994].

The pay-off from investing in good persistent technology is only apparent when a large application is used over a long period. Hence it is difficult to properly evaluate this technology without further investment in building up a realistic load, writing a reasonable volume of application software and conducting the evaluation over a long period [Atkinson, 1992, Atkinson, et al., 1993]. Currently, none of the systems that build up persistence from the operating system kernel have received enough developmental effort to enable them to support such experimentation.

4.2 Implementing Persistence by Reachability

Support systems conforming to the last two architectures and many that are essentially extended systems implement:

- persistence by reachability;
- automated data movement; and
- re-use of space.

The basic algorithms for these three crucial tasks, which are applicable in all three architectures, are now presented. This is followed by a brief introduction to some of the ways in which the algorithms may be refined.

4.2.1 Incremental Loading and Swizzling

The persistent system will provide some means of naming and loading persistent roots — in the kernel one root will suffice, since other naming schemes and multiple roots can be implemented with persistent data structures and code accessed by this primitive root. The persistent root will contain references to other objects which will themselves contain references to further objects. Data movement into the active store (that against which code operates) is initiated when such references are de-referenced in an attempt to use the object to which they refer. These objects on which transfers are based may not correspond exactly to those which the programmers think about, as compilers may introduce mappings.

When a value is accessed, the support system arranges that it is loaded before computation is allowed to continue. Typically, there are two forms of address that refer to a value:

- a Persistent Identifier (PID) that has long-term reliability — e.g. it is some logical identification of the object, interpreted through a mapping or index, or (with less flexibility and permanence) it is a machine, device and disk address; and
- a Local Address (LA) which is the form used by the executing process on the current hardware — e.g. a virtual memory address.

Different stores utilise different means of discriminating these two forms and some operate with only one performing both roles, these issues are discussed later. It is desirable to arrange that if a LA is encountered, the evaluation is not interrupted. However, if a PID is de-referenced (an attempt is made to use the value to which it refers), then the action in Table 5 is taken.

Table 5: Action Taken on De-referencing a PID

1	discover whether the referend is already resident
2	if it is (false object-fault), overwrite the PID with the corresponding LA (swizzle [Moss, 1990, Wilson, 1990]), so that future de-references occur at normal computational speed
3	if it isn't (object-fault), read in the object, and keep a record of this PID to LA mapping to support steps 1 and 2, and then swizzle, i.e. overwrite the PID with the LA.

This algorithm will arrange that any long-lived data required by an application is automatically made available to that program. Note that in step 3 of the above algorithm, it is necessary to record the relationship between a PID and its current local address in this process. This information is stored in a two-way accessible table, called here the PIDLAM (PID to Local Address Map). It is called the “Resident Object Table” in [Kemper & Kossmann, 1995]. It has to be accessed by PID in step 1 of the above algorithm to preserve sharing of substructures and by LA during de-swizzling (see 4.2.2).

Figure 12: Lazy Incremental Loading Algorithm

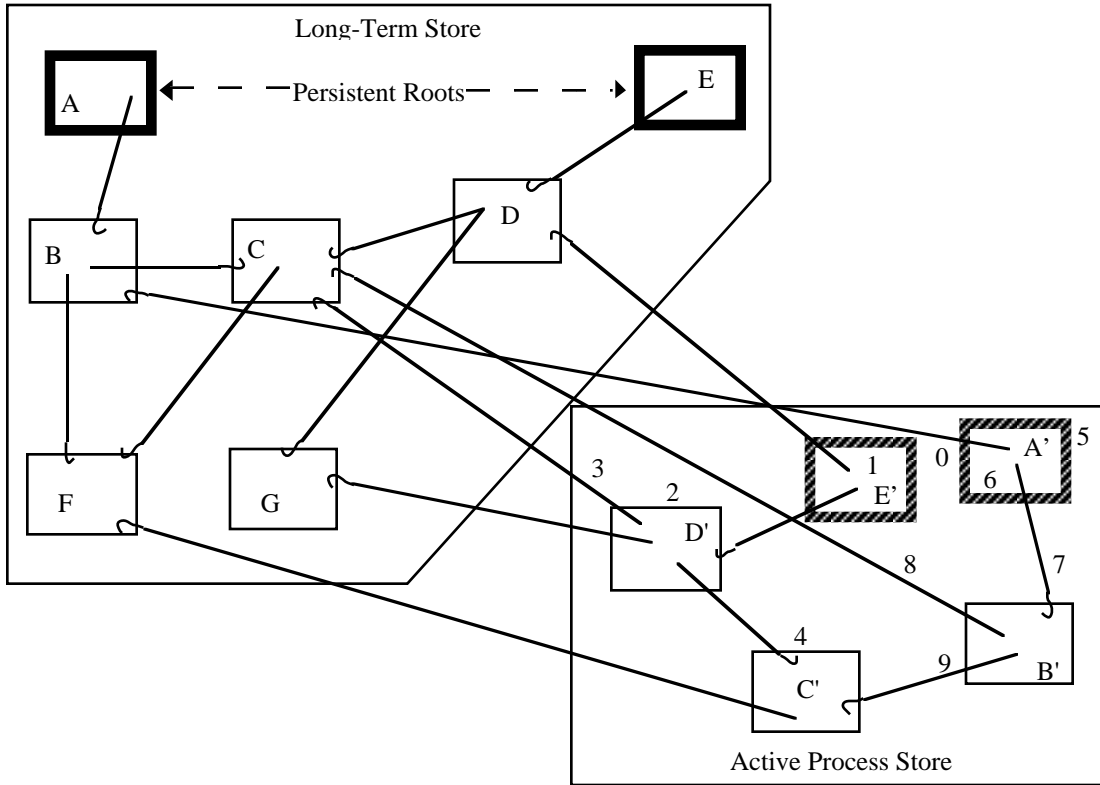


Figure 12 is used to illustrate this automatic incremental loading process. It shows an intermediate store state during the execution of a persistent program. The boxes represent objects which may be in the long-term or short-term stores or in both. The lines denote references with the hook at the end corresponding to the referend. Thus, if we consider them attached to the object that holds them, the persistent objects are all those obtained by starting at each persistent object and pulling on all the hooks that emanate from those objects.

The relevant steps in the program execution that leads to the store state shown in Figure 12 are given in Table 6.

Table 6: Steps in Incrementally Loading Data in Figure 12

0	Obtain persistent root E as E'
1	Attempt to de-reference PID to D — object-fault
2	Copy D as D' into active store and Swizzle: replace PID to D with LA of D'
3	Attempt to de-reference PID in D' to C — object-fault
4	Copy C as C' into active store and Swizzle: replace PID to C with LA of C'
5	Obtain persistent root A as A'
6	Attempt to de-reference PID in A' to B — object-fault
7	Copy B as B' into the active store and Swizzle: PID to B with LA of B'
8	Attempt to de-reference PID in B' to C — false object-fault
9	Find PID of C from 4 and Swizzle: PID to C with LA of C'

4.2.2 Promotion to Persistence on Checkpoint

Data flow from the active store to the long-term store also has to be automated. Application programs from time to time will request checkpoint, often as part of commit. At this point any objects that have been brought from long-term store and

then changed must be written back. In addition, new objects that are now reachable from the persistent root(s) must now be written to the long-term store, after they have been allocated space and a PID. As each object is written out, all of the LAs in the object must be de-swizzled, so that they now refer to PIDs. PIDs have to be allocated in advance of the transfers to correctly implement common sub-structures and cycles.

It should be noted that an object can only be promoted to longevity if some already persistent object has been updated to hold a reference to it, or if it is referenced by an object that itself is being promoted to longevity. The promotion algorithm is shown as Table 7.

Table 7: Algorithm to Promote Objects to Longevity

1	$Mutated := \{ \text{all the mutated persistent objects} \}$
2	$Promotions := \{ \text{every new object directly reachable from } Mutated \}$
3	$Promotions := Promotions \mathbf{union} \{ \text{every new object directly reachable from } Promotions \}$
4	repeat step 3 until no more additions
5	allocate PIDs for all objects in $Promotions$
6	transfer all members of $Mutated \mathbf{union} Promotions$ to long-term storage

The effect of this algorithm is illustrated with the aid of Figure 13.

Figure 13: Part of the Active Store after Mutation

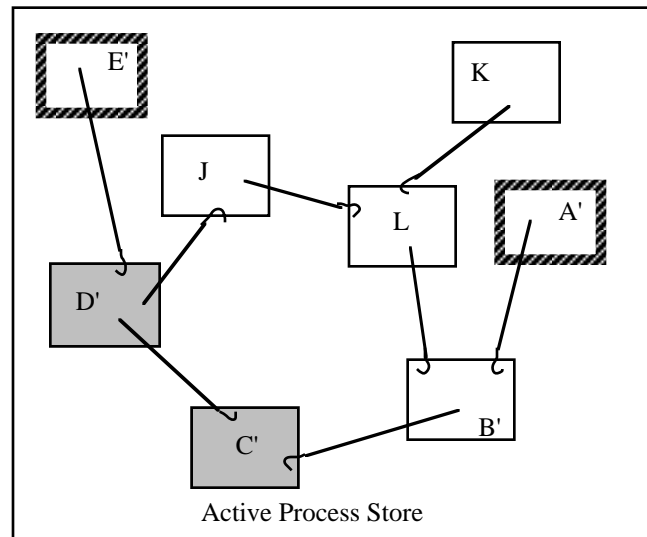


Figure 13 shows part of the active store after a computation has continued from the incremental loading shown in Figure 12. It has mutated D' and C' , and has created three new (transient) objects J , K and L . It has overwritten the reference in D' to G with a reference to J which contains a reference to L which contains a reference to B' . The object K references L but is not referenced by any persistent or newly formed object.

The steps in the checkpoint on the store in this state are shown in Table 8.

Table 8: Checkpoint of Store in Figure 13

1	Discover <i>Mutated</i> = { D', C' }
2	Set <i>Promotions</i> = { J }
3	Set <i>Promotions</i> = { J, L }
4	Allocate PIDs for J and L
5	De-swizzle pointers in D', C', J and L
6	Copy D', C', J and L to long-term storage
7	If processing is to continue, swizzle pointers in D', C', J and L

Carrying out these steps both efficiently and precisely without unduly slowing the execution before the checkpoint requires careful engineering (see below). The copying in step 6 is usually performed in such a way that recovery is possible after a machine failure during checkpoint. For example, logs are written [Agrawal & De Witt, 1985, Moss & Sinofsky, 1988, Ruffin, 1992, Scheuerl *et al.*, 1995] or shadow pages are used [Brown, 1987, Munro, *et al.*, 1994].

4.2.3 Garbage Collection and Object Termination

Persistence based on reachability implies that when an object becomes unreachable it is no longer useful — no future program execution can access it. However, it will still occupy space in the long-term store and an opportunity exists to reclaim that space. An obvious method, also based on reachability, is garbage collection.

Since the long-term store is usually held on devices such as disk and may grow very large, these garbage collections may be infrequent and incremental. The implementor trades space and perhaps density of useful data against garbage collection costs.

Garbage collection is traditionally used solely to recover space [Wilson, 1992]. Garbage collections of the active store interact with the preceding algorithms in several important ways.

- When they fail to recover sufficient space by normal methods, they can release space by discarding unmodified objects that were previously loaded — this requires that pointers that once referenced these discarded objects are de-swizzled (a major issue in [Kemper and Kossmann, 1995]).
- When, even after discards of unmodified objects, there is still insufficient space, they can copy out modified objects or prematurely (and possibly falsely) promote new objects (steps 4 to 7 of the checkpoint algorithm must be applied) — this increases the complexity of the subsequent checkpoint algorithms and is best avoided if at all possible.
- They can incrementally copy data to log files for recovery purposes [Kolodner, 1987, Kolodner *et al.*, 1989].

These additional operations may also be invoked during garbage collection because the density of objects still in use has fallen to a level where thrashing has set in [Fenichel & Yochelson, 1969]. Alternatively, dynamic clustering may be used to reduce thrashing [Benzaken *et al.*, 1991].

4.2.4 Engineering Issues

Implementing the preceding algorithms with the support of conventional operating systems and hardware presents several challenges for store designers. A few of the choices are illustrated here.

The long-term form of the data may be exactly the same as that required for the active store. This avoids the costs and complexity of translation, swizzling and maintaining a PIDLAM. It has the disadvantage of not scaling very well and of not guaranteeing

that the data (which normally includes code) can be moved to a new platform as hardware architectures change [Atkinson, et al., 1993].

If the same representation prevails, then for small, single-user systems a total copy into virtual memory will be fast and adequate. For larger and shared systems of this kind, memory mapping techniques can be used [Koch *et al.*, 1990, Rosenberg *et al.*, 1990, Brown, et al., 1992, Singhal, et al., 1992, Munro, et al., 1994].

Object-faults (and false object-faults) can be detected by using the address translation unit's (ATU) protection or by in-line code before each de-reference. The former normally has the overhead of several transitions across the operating system boundary and the later has costs even after the objects have been loaded.

Data movement can be on parts of objects, on objects or on physical divisions (e.g. pages or groups of disk blocks) that match the hardware. The first has the advantage of accommodating very large objects, but has the disadvantage of increasing the complexity of object-fault detection (e.g. the next part of an image may need loading as a raster operation traverses it). The first and second have the advantage of not cluttering the working set with co-resident objects, but they require intermediate buffering and hence additional copying (Kemper and Kossmann show this has performance benefits [Kemper and Kossmann, 1995]). These two methods avoid phantom locking and checkpointing of co-resident objects, but complicate the implementation of resilience. The third method has the advantage that the actual transfer exploits the disk and channel well, and may even gain from optimisations made to support paging. It suffers because it brings into active memory co-resident objects that are not used. This detrimental effect will increase as a PAS ages unless a good dynamic clustering algorithm is operational.

Swizzling can be carried out at various times. Swizzling every time a PID is dereferenced has the advantage of avoiding the need to overwrite and later de-swizzle, but costs a look-up in the PIDLAM at every de-reference. This might be feasible if there was associative hardware supporting the PIDLAM [Russell, 1994].

Swizzling only on the first de-reference of a PID has the advantage of speed where algorithms repeatedly traverse the objects once they are loaded, but it involves updates to objects which may interfere with the use of hardware protection to discover the *Mutated* set or to record references to more recent generations (in generational garbage collectors [Cook *et al.*, 1993]).

As an object is loaded all its PIDs can be swizzled. Wilson uses memory mapping to do this a whole page at a time [Wilson, 1990]. Wilson's scheme has the advantage of avoiding the costs of detecting any false object-faults, and it avoids system initiated writes to objects. It has the cost of allocating virtual address space more rapidly and of restricting the variation between PID and LA to the page number part of the address, thus preventing objects from expanding as they are translated.

Translation of objects permits much longer longevity — they can be held in a canonical form that does not derive from a particular architecture. This translation may be combined with decompression (on load) and compression on transfer to longer-term storage. This can be performed on a page or on an object at a time.

Discovering the *Mutated* set at checkpoint may be achieved by generating code during compilation that will insert objects into the *Mutated* set or by using the ATU's protection system to detect an update dynamically. The former has the disadvantage that it incurs execution penalties on subsequent updates of the object (unless the compiler did some clever optimisation) as in [Moss & Hosking, 1994]. The latter has the disadvantage that it is only approximate (there are likely to be other objects in the same protected region) and interacts with other uses of the ATU's protection and the system initiated updates.

The above list is not exhaustive. Other major issues include whether to allow objects to span pages, whether to cluster and prefetch, where to allocate space and PIDs on

promotion, whether to use shadow paging or logging, should a PID encode a store location, etc. Designing a good persistent object store is still a considerable challenge, as the interaction of this plethora of choices is only understood where it has been sampled.

4.3 Type Safe Linguistic Reflection

Type safe linguistic reflection is defined in [Stemple, et al., 1992b] as the ability of a running program to generate new program fragments and to integrate these into its own execution. This is the basis for system evolution which itself is necessary to achieve adequate PAS longevity. For safety reasons only strongly typed reflection will be considered.

Linguistic reflection has the goal of allowing a program's behaviour to adjust dynamically in order to provide flexibility and high productivity. It thus extends the data modelling of the type system and it should not be surprising therefore to find a tension between type systems and reflection. The possibility that a program may significantly change its behaviour decreases the opportunity for static type checking and thus compromises some of the benefits of typing. Thus the reflective facilities are controlled in a manner designed to retain as much static type checking as possible without the control being so severe as to remove all the benefits.

The two techniques for type-safe linguistic reflection that have evolved are: compile-time linguistic reflection and run-time linguistic reflection. Compile-time linguistic reflection [Stemple *et al.*, 1990, Stemple *et al.*, 1992a] allows the user to define generators which produce representations of program fragments. The generators are executed as part of the compilation process. Their results are then viewed as program fragments, type checked and made part of the program being compiled.

Run-time linguistic reflection [Dearle and Brown, 1988, Kirby, 1992a, Kirby *et al.*, 1994] is concerned with the construction and binding of new components with existing components in an environment. The technique involves the use of a compiler that can be called dynamically to compile newly generated program fragments, and a linking mechanism to bind these new program fragments into the running program. Type checking occurs in both compilation and binding.

The benefits of type safe linguistic reflection in database and persistent programming consist mainly of two capabilities. The first is the ability to implement highly abstract specifications, such as those used in query languages and data models, within a strongly typed programming language. The second is the ability to accommodate some of the continual changes in data-intensive applications without resorting to *ad hoc* restructuring methods. Both capabilities involve reflective access to the types of a system that is changing itself and both approaches avoid introducing extra levels of interpretation.

Both compile-time and run-time reflection have been provided in previous languages. Compile-time reflection appears in the macro facilities of Scheme [Rees & Clinger, 1986] and POP-2 [Burstall *et al.*, 1971]. Run-time reflection appears in the *eval* functions of Lisp [McCarthy, et al., 1962] and SNOBOL4 [Griswold *et al.*, 1971] and the *popval* function of POP-2 [Burstall, et al., 1971].

Type safe linguistic reflection is different for the following reasons.

- More information is available to the reflective computation, in the form of systematically acquired types. This information can be used to automatically adjust to implementation details and system evolution. Linguistic reflection without strong typing has little systematic information available about the structures involved in computation.
- The type safety of all newly generated program fragments is checked before they are allowed to be executed. Such type discipline is highly advantageous in a database programming environment in which the integrity of long-lived data must be maintained.

It is somewhat ironic that strong typing, which makes it difficult to integrate reflection with typed programming languages, is what makes linguistic reflection effective as an amplifier of productivity.

Type safe linguistic reflection has been used to attain high levels of genericity [Stemple, et al., 1990, Sheard, 1991] and accommodate changes in systems [Dearle and Brown, 1988, Dearle *et al.*, 1990]; two examples of these are given below. It has also been used to implement data models [Cooper, 1990a, Cooper, 1990b, Cooper & Qin, 1992], optimise implementations [Cooper *et al.*, 1987, Fegaras & Stemple, 1991, Cutts *et al.*, 1994] and validate specifications [Fegaras *et al.*, 1992, Stemple, et al., 1992a]. The importance of the technique is that it provides a uniform mechanism for software production and evolution. A formal description of linguistic reflection is given in [Stemple, et al., 1992b], here we concentrate on the uses.

4.3.1 Uses of Type Safe Linguistic Reflection

Two examples of type safe linguistic reflection are presented, in order to show how the reflection mechanisms appear at the level of a programming language. The examples are abstraction over types and accommodating evolution in strongly typed persistent systems.

4.3.1.1 Abstraction Over Types

As indicated earlier a generic natural join function provides an example of abstraction over types that is beyond the capabilities of most polymorphic type systems. The details of the input types, particularly the names of the tuple components, significantly affect the algorithm and the output type of the function, determining:

- the result type;
- the code to test whether tuples from both input relations match on the overlapping fields; and
- the code to build a relation having tuples with the aggregation of fields from both input relations but with only one copy of the overlapping fields.

The type of a polymorphic natural join function would be

$$\forall a. \forall b. \forall c. (\text{set } [a] \times \text{set } [b] \rightarrow \text{set } [c])$$

That is the function takes as parameters two sets and returns a third set as a result. This function cannot be written as a statically-typed polymorphic procedure since it requires knowledge of the type structure of a , b and c and they are explicitly abstracted over here. However the type structure is known for any particular call. Linguistic reflection allows a metafunction to be written, the generic natural join, that can interrogate the structure of the input types. At the point of use of the natural join, the metafunction is called by the compiler and given as input the type structure. This compile-time invocation will compute the specific natural join function for these input types and a call to the specific natural join. All of this generated code is then fed into the compiler for checking and code generation.

The writing of the metafunction involves computing over the representations of the input types. Thus the language must provide facilities for this. Indeed it is this ability that contrasts the reflective language with other polymorphic systems. When using run-time reflection it is also possible to examine the size and representation of the input sets and to generate an optimised algorithm.

4.3.1.2 Evolution in Strongly Typed Persistent Systems

Type safe linguistic reflection may also be used to accommodate evolution within strongly typed persistent object stores. A characteristic of such stores is that the set of types of existing values in the store evolves independently from any one program. This means that when a program is written or generated some of the values that it may have to manipulate may not yet exist, and their types may not yet be known for inclusion in the program text.

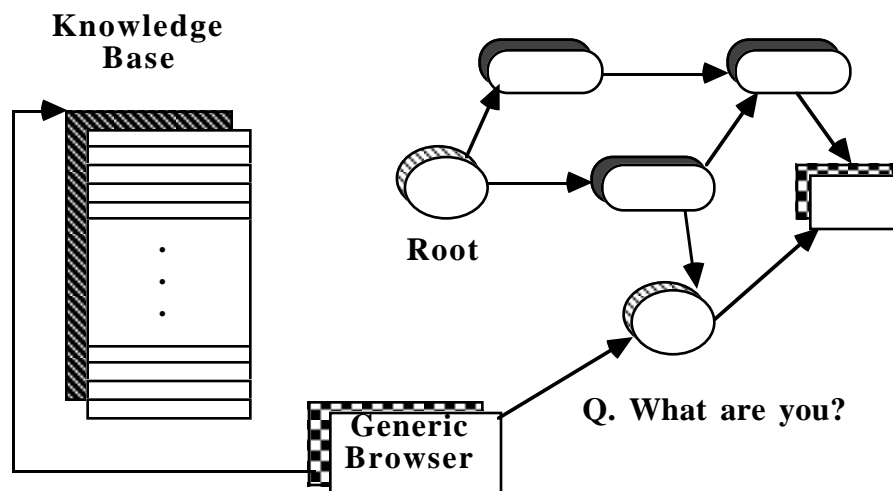
An example of such a program is a persistent object store browser [Dearle and Brown, 1988] which displays a graphical representation of any value presented to it. The browser may encounter values in the persistent store for which it does not have a static type description. This may occur, for example, for values that are added to the store after the time of definition of the browser program. For the program to be able to denote such values, they must belong to an infinite union type, such as type **any** described earlier.

However, the program cannot contain static type assertions for all the types that may be encountered as their number is unbounded. There are two possibilities for the construction of such a program: it may either be written in a lower-level technology using interpretation [Kirby & Dearle, 1990] or else be written using linguistic reflection.

To allow a reflective solution the program must be able to discover dynamically the specific type of a value of the union type. Such functionality may be provided in a strongly typed language, without compromising type security, by defining representations of types within the value space of the language.

Figure 14 illustrates the use of linguistic reflection to define programs that operate over values whose type is not known in advance. The generic browser takes a specification of the object type and generates a program to browse over it. It may also store that program in the knowledge base for future use should it encounter an object of the same type again. These programs potentially perform different operations according to the type of their operands but without endangering the type security of the system or invoking an extra layer of interpretation. The requirement for such programs is typical of an evolving system where new values and types must be incrementally created without the necessity to re-define or re-compile existing programs.

Figure 14: A Persistent Store Browser



5 Achievements

The progress of persistence research can be charted by examining the proceedings of the regular workshops held by its practitioners. There are two series of international workshops: the POS — on persistent object systems — and the DBPL — on database programming languages — and several *ad hoc* events such as persistence tracks at HICSS — the Hawaii International Conference on System Sciences. These are tabulated in Table 9 for those who wish to scan this literature.

Table 9: Major International Meetings of Persistence Researchers

	Date	Venue	Organisers
POS 1	Aug. 1985	Appin, Scotland	Atkinson, Buneman and Morrison [Atkinson, et al., 1988b]
POS 2	Aug. 1987	Appin, Scotland	Atkinson and Morrison [Atkinson <i>et al.</i> , 1987]
DBPL 1	Sept. 1987	Roscoff, France	Bancilhon and Buneman [Bancilhon & Buneman, 1990]
HICSS 22	Jan. 1989	Hawaii, USA	Atkinson and Morrison [Atkinson & Morrison, 1989]
POS 3	Jan. 1989	Newcastle, Australia	Koch and Rosenberg [Rosenberg and Koch 1989]
DBPL 2	June 1989	Salishan, Oregon, USA	Hull, Morrison and Stemple [Hull, et al. 1989]
Security	May 1990	Bremen, Germany	Rosenberg and Keedy [Rosenberg & Keedy, 1990]
POS 4	Sept. 1990	Martha's Vineyard, USA	Dearle, Mitchell and Zdonik [Dearle <i>et al.</i> , 1990]
DBPL 3	Aug. 1991	Nafplion, Greece	Kanellakis and Schmidt [Kanellakis & Schmidt, 1991]
HICSS 25	Jan. 1992	Hawaii, USA	Morrison and Atkinson [Morrison & Atkinson, 1992]
POS 5	Sept. 1992	San Miniato (Pisa), Italy	Albano and Morrison [Albano & Morrison, 1992]
DBPL 4	Aug. - Sept. '93	Manhattan, NY, USA	Beeri, Ohori and Shasha [Beeri, et al. 1993]
POS 6	Sept. 1994	Tarascon, France	Atkinson, Benzaken and Maier [Atkinson <i>et al.</i> , 1995]
HICSS 28	Jan. 1995	Hawaii, USA	Rosenberg and Dearle [Rosenberg & Dearle, 1995]

From the contents of the proceedings cited in Table 9 it is apparent that a large community of researchers are addressing persistence. This section reports some of their achievements. The persistent languages that have been developed are displayed in Table 10.

Table 15: Milestones in Persistent and Database Programming Languages

1977	Pascal/R [Schmidt, 1977]
1980	SmallTalk [Goldberg & Robson, 1983], PS-algol [Atkinson, et al., 1983], Plain [Wasserman <i>et al.</i> , 1981], Taxis [Mylopoulos <i>et al.</i> , 1980]
1981	Daplex [Shipman, 1981], Adaplex [Smith <i>et al.</i> , 1981, Chan <i>et al.</i> , 1987]
1983	Galileo [Albano, et al., 1985], Modula/R [Koch <i>et al.</i> , 1983], Persistent procedures [Atkinson and Morrison, 1985]
1984	Amber [Cardelli, 1986], Persistent Prolog [Bocca & Bailey, 1987]
1985	CPS-algol experiment [Krablin, 1987b], Poly [Matthews, 1982], OPAL [ServioLogic Ltd, 1987]
1986	DBPL [Matthes & Schmidt, 1989], RAPP [Hughes and Connolly, 1990]
1987	Quest [Cardelli, 1989], E [Richardson and Carey, 1989], X [Hurst & Sajeev, 1990]
1988	DPS-algol [Wai, 1989], Napier88 [Morrison, et al., 1989]
1991	P-Quest [Matthes <i>et al.</i> , 1992], Staple [Davie & McNally, 1990], P-Galileo [Brown, et al., 1992], O ₂ [Deux, 1990]
1992	Hyper-programming [Kirby, et al., 1992], Commercial Persistent C ⁺⁺
1993	Tycoon [Matthes & Müßig, 1993], Fibonacci [Albano, et al., 1995]
1994	Napier88 version 2 [Morrison, et al., 1994]

A few aspects of the use of these languages are presented below.

5.1 Incremental Construction

Because PASs are long-lived their construction usually takes place incrementally. As each requirement is recognised or as resources become available, new suites of software and collections of data are added to the operational system. Databases have supported well the incremental addition of programs to the operational PAS via their dynamic binding to schemata or views. They have made less provision for the incremental growth of schemata taking the unrealistic view that the entire schema can be designed at the outset and that only minor modifications need be made using a schema editor. In reality, use of the schema editor is frequent, but there is little support for change in data models.

The persistence technology so far developed makes adequate provision for the incremental addition of collections of programs by providing dynamic binding constructs [Atkinson *et al.*, 1988a]. This is widely exploited in the prevalent persistent programming style [Dearle, 1988, Connor, 1990, Cutts, 1992, Sjøberg, 1993]. Incremental addition and replacement of program parts tends to use smaller incremental units than are used with the traditional database technology. This is because the retention of types and structure and the provision of persistent bindings reduces the complexity and cost of using smaller units.

Persistence technology takes a different view over the construction of schema. It is normal to consider the design of types at the same time as the programs that will use them. The technology itself therefore, has no bias towards programs either being more or less incremental than types. This freedom from bias appears to have pay-offs. The outline design of the data can be completed early and stored as a set of type definitions, but much of the detail can be postponed using the dynamic checking points. Then, as programs are required, they and their associated detailed types are designed together. This has the advantage of requiring less perturbation of the existing types and of allowing designers and programmers to concentrate on a particular subsystem's types. To facilitate this, persistent systems have developed mechanisms for collecting types together to record and hopefully preserve this

modularity, for example, Quest has modules [Cardelli, 1989] and Napier88 declaration sets [Kirby, et al., 1994].

5.2 Hyper-programming

As mentioned earlier, the presence of a persistent environment in which programs are composed, compiled, linked and run extends the scoping of bindings to persistent values. This means that objects accessed by a program may already be available at the time that the program is composed. Where the source programs are themselves values, static R and L-value bindings may be made in the program source to persistent values. Thus these bindings, called links in this case, can be included in the program instead of the more traditional textual descriptions of where to find persistent values. The source text of the program now contains some text and some links to persistent values and as such it is a non-flat representation of the program. By analogy with hyper-text, a program containing both text and links to persistent values is called a *hyper-program* [Kirby, et al., 1992].

Figure 16, taken from [Kirby, et al., 1992], shows an example of a hyper-program. The links embedded in it are represented by non-textual tokens to allow them to be distinguished from the surrounding text. The first link is to a first-class procedure value *writeString* which writes a prompt to the user. The program then calls another procedure *readString* to read in a name, and then finds an address corresponding to that name. This is done by calling a procedure *lookup* to look up the address in a table data structure linked into the hyper-program. The address is then written out. Note that code objects (*readString*, *writeString* and *lookup*) are denoted using exactly the same mechanism as data objects (the table). Note also that the object names used in this description have been associated with the objects for clarity only, and are not part of the semantics of the hyper-program.

Figure 16 : A Hyper-program

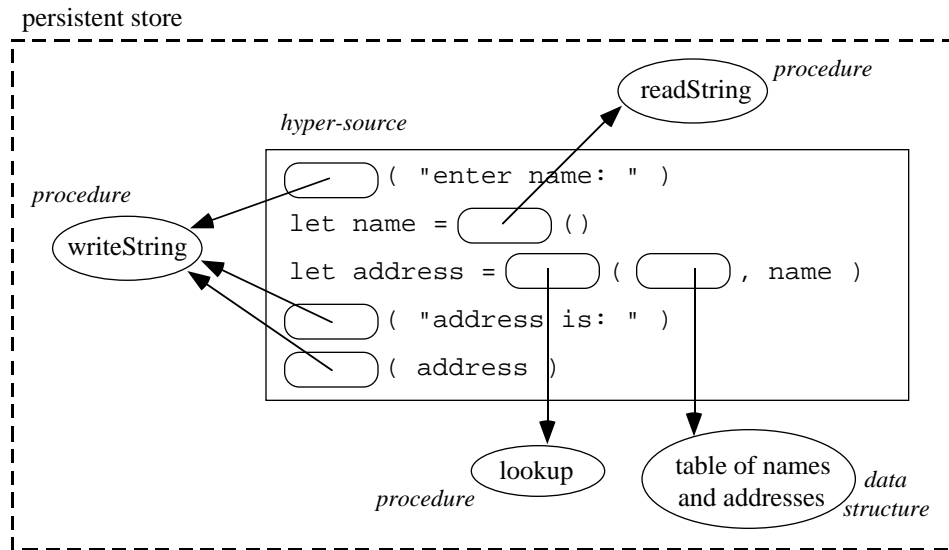
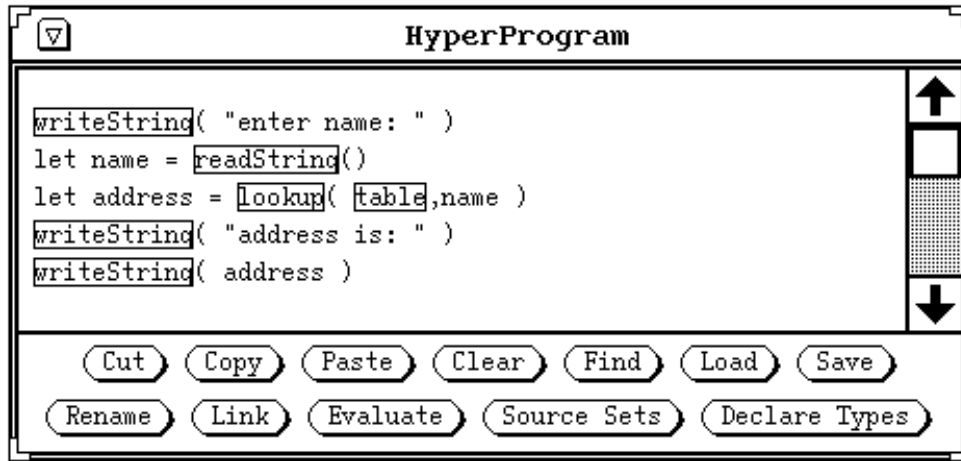


Figure 17, again taken from [Kirby, et al., 1992], shows an example of the user interface which might be presented to the programmer by a hyper-program editing tool. The editor contains embedded light-buttons representing the hyper-program links; when a button is pressed the corresponding object is displayed in a browser window. The browser is also used to select persistent objects for linking into hyper-programs under construction.

Figure 17 : User Interface to a Hyper-program Editor



The benefits of hyper-programming are discussed in [Farkas *et al.*, 1992, Kirby, 1992b, Kirby, et al., 1992] and include:

- being able to perform program checking early — access path checking and type checking for linked components may be performed during program construction;
- being able to enforce associations from executable programs to source programs — links between source and compiled versions may be used;
- support for source representations of all procedure closures—free variables in closures may be represented by links, thus allowing hyper-programs to be used for both source and run-time representations of programs; and
- increased program succinctness—access path information, specifying how a component is located in the environment, may be elided.

5.3 Persistent Workshop and Methodologies

Persistent programmers need virtually all the tools in a programming environment that are required by other programmers (editors, compilers, build-managers, version-managers, etc.). There are, however, some important differences in their requirements and by exploiting persistent technology to provide this environment, their needs can be met in novel ways.

An experiment is underway in building such an environment called the persistent workshop. The experiment has three purposes:

- to demonstrate the ease of construction and new functionality afforded by developing the workshop using orthogonal persistence;
- to provide a good programming environment for persistent programmers; and
- to set up an example of an operational PAS so that its usage and behaviour may be studied.

All the source and operational code and any design material are held in the persistent store. All the tools communicate via the persistent store. Dynamically typed interfaces are used where independent tool evolution requires flexibility. The workshop provides a set of tools and an extensive library of components [Atkinson, et al., 1993, Kirby, et al., 1994]. Examples of subtly changed requirements on these tools arising from persistence include:

- tools that provide a means of configuring a workbench, establishing a style of working and a particular view of the persistent store — these enable programmers to limit the complexity of data and choices open to them to those relevant to their task [Waite, 1995];

- tools to examine the relationship between the persistent store and programs — these enable programmers to more quickly appreciate the existing structures and to verify the results of their work [Lavery, 1995a, Lavery, 1995b];
- tools to organise and build libraries of program parts and data, and to search these libraries using information retrieval techniques — which are intended to improve component re-use both within a PAS and among PASs [Brown, 1993];
- tools to manage the large number of incremental units, including providing aids for building and installing subsystems and to verify that the parts comply with a methodology [Sjøberg *et al.*, 1995] and
- tools to assist with change management [Sjøberg, 1991].

These tools already exploit properties of persistence. For example, they leave data structures in the persistent store to accelerate subsequent operations. As another example, the build manager, the store analysers, the methodology checker and the change management aids all rely on being able to scan the store and on knowing that, unlike many systems, all external information, vital for its construction, is explicitly connected to the PAS's representation.

Two methodologies are supported in the workshop at present. One is closely allied to traditional database practice. In that, the PAS construction proceeds via the design of a system in some conventional data model [Cooper & Qin, 1994]. The types and programs are then automatically derived. This works reasonably well on a small scale, but will not yet scale up as incremental construction via automated generation is still a research issue. The other uses a methodology called SPASM [Sjøberg, et al., 1995] and sets out to delimit dependencies between incremental components accurately. This involves the choice of particular ways of using the persistent technology but has the advantage that more of the build and change processes can be automated.

5.4 Persistent Software Engineering

The concept of the hyper-program, containing links to persistent objects and made possible by the provision of strong typing and referential integrity has a number of applications in software engineering environments. Three examples of such usage are: the Napier88 system, the Flex system [Currie, 1985, Stanley, 1986, Stanley and Drummond, 1988] and the Vesta Configuration Management System [Chiu & Levin, 1993, Levin & McJones, 1993].

The Napier88 system contains a fully functional hyper-programming system and in addition contains prototype versions of the following further uses of hyper-links in the software engineering context [Morrison, et al., 1995]:

- simplification of the programming model via hyper-code;
- version control;
- configuration management; and
- documentation.

One of the advantages of hyper-programming is the ability to use the hyper-program representation for both source and run-time representations of programs. At program composition time, the programmer may construct a hyper-program using a tool which is a combination of an editor and a browser. At run-time the hyper-program representation may also be used to represent an active computation. This is possible due to the non-flat nature of that representation. Free values, that is, non-local references to objects and procedures, may be represented as links and the inherent sharing of values and locations referred to by links is preserved. This is not possible with textual representations of programs since the sharing is lost.

The hyper-code abstraction allows a single program representation, the hyper-program, to be presented to the programmer at all stages of the software development process. In constructing a program, the programmer writes hyper-code. During

execution, during debugging, when a run time error occurs or when browsing existing programs, the programmer is presented with, and only sees, the hyper-code representation. Thus the programmer need never know about those entities that the system may support for reasons of efficiency, such as object code, executable code, compilers and linkers. These are maintained and used by the underlying system but are merely artefacts of how the program is stored and executed, and as such are completely hidden from the programmer. Only one set of tools is required for manipulating the hyper-code thereby simplifying the user interface and the complexity of the system implementation. This permits concentration on the inherent complexity of the application rather than on that of the support system.

Software environments are made safer but less flexible whenever a name binding can be replaced by a hyper-link. The replacement of name bindings found in traditional software systems by links and reverse links ensures that a given software component and the objects to which it links remain accessible from one another, as a consequence of persistence defined by reachability. Thus in version control, configuration management and documentation systems, the software items can refer directly to the components to which they relate. Program and documentation may be directly linked together with hyper-links, a particular configuration may have a hyper-link to a configuration script and a version may have direct links to related versions.

By using similar methods, the same advantages accrue to other activities supported by software environments which are not discussed here, such as debugging, profiling and optimisation [Cutts, 1992].

The Flex system [Currie, 1985, Stanley, 1986, Stanley and Drummond, 1988] consists of a programming language supported by a persistent file store that contains structured data. Since the Flex language is a complete version of algol-68 it has all the ingredients required for the exploitation of strongly typed persistent linkage, namely: persistent links (capabilities in this case), higher-order functions and strong typing. Flex also has the concept of hyper-links, called cartouches after their user interface representation. Source code may contain cartouches which point to persistent objects, in this case typed files. The notion of pairing source code with executable code so that they can only be updated in lock-step is also present in Flex.

The Vesta configuration management system [Chiu and Levin, 1993, Levin and McJones, 1993] is based on a file system that does not allow overwriting in place. Since the system is Unix based, this means that there is no write permission on any file and that inodes are unique. By this means, configurations are guaranteed to refer to the same file every time since the file system ensures the integrity of the inode. To support this, the implementors re-wrote the Unix file system to remove update in place and thereby support configuration management by persistent links with referential integrity.

6 Future Directions

An overview of the current state of persistent system research has been presented. There are commercial applications [Greenwood *et al.*, 1992] and many experimental ones [Grossman *et al.*, 1994, Reinwald, et al., 1994]. These indicate the promise of orthogonally persistent object systems and that the technology is already usable. However, much remains to be done to achieve the full potential of persistence and to establish this approach as one which can be safely adopted by industry. These further avenues of research and development are presented under three headings:

- **Extensions of Persistence.** Developments that will make the facilities integrated within orthogonally persistent systems more complete;
- **Exploitation of Persistence.** Persistent methodologies, evaluation and use of the existing persistent technology; and
- **Delivering Persistence.** Continuing improvement to the engineering of persistent support systems.

6.1 Extension of Persistence

Extensions of persistence fall into two categories: those which are necessary to facilitate PAS design, construction, maintenance and operation already well established in the database context and those which meet new needs currently unmet by any technology.

To meet the first requirement, which is here called “database completeness” the facilities for data modelling, bulk types, transactions, distribution, autonomy and evolution all deserve further attention. The second requirement overlaps with the first, in that better data models informed by research into types, more varieties of transactions, complex transactions that integrate well with recovery, etc. are required in both contexts. New extensions in the second category include seeking the conceptual simplicity of persistence in the combination of programming languages with other sub-systems, such as: user-interface management systems, operating systems and communication systems.

6.1.1 Conceptual Modelling and Generic Interfaces

The use of types to describe data was illustrated earlier in the paper. Some would argue, however, that these type notations are insufficiently descriptive or fail to exploit the body of knowledge regarding database design using: E-R models, relational models, semantic models, object modelling, etc. There will always remain application domains where a new group of data modelling constructs will be useful. The utility of providing extensible libraries of type constructors requires validation in a realistic PAS development environment.

New type constructors can be defined and *named* using the pre-existing or pre-defined constructors. Cooper has extensively explored the process of mapping traditional data models to types in several different languages and has used reflection to build aids to performing this mapping automatically [Cooper, 1990a, Cooper and Qin, 1992, Cooper and Qin, 1994]. His systems automatically generate all the types that match some data model and both a text-based and diagram-based schema editor. They also generate the usual interface components (e.g. forms generators, etc.), skeleton application programs and a relevant library of persistent components. These should aid the transfer of existing system design skills to the new technology, migration of designs, rapid prototyping and application construction. That expectation needs to be validated via real applications.

Others have developed more specific mappings from models to persistent type systems [Stemple, et al., 1992a, Sheard & Hook, 1994, Wetzal, 1994, Albano, et al., 1995]. It is likely that this approach of mapping to type systems using reflection will prove widely applicable and will relieve the pressure on type systems to develop to the levels of complexity where they are difficult to understand and expensive to implement.

It remains to be seen whether the code and constructs to accommodate changes specified in terms of such higher-level models can be automatically generated and applied. If this is achieved then the provision of high levels of data independence can be automated for a wide range of modelling styles [Atkinson, et al., 1993].

It is probable that there are considerable benefits to be gained by generating code from many high-level notations into a common orthogonally persistent system. The composite would have consistent behaviour under stress (e.g. recovery). It is expected that the use of reflective generation will allow such a persistent target language to be simplified. The authors do not know of any exploration of this potential.

6.1.2 Bulk Types, Query Languages and Optimisation

An important component of data models is bulk types [Atkinson and Buneman, 1987, Atkinson *et al.*, 1991]. Their value is three-fold:

- they abstract over the size of collections they represent;

- they allow the description of important regularities in the data; and
- they permit exploitation of the regularity, e.g. in query notations and in optimisation of operations and data movement.

Recent work shows that certain classes of bulk types can be treated consistently [Ohori, et al., 1989, Trinder & Wadler, 1989, Breazu-Tannen *et al.*, 1991, Buneman *et al.*, 1994]. It has also been shown that these bulk types can be implemented by mapping them to the underlying system [Cooper and Qin, 1992, Ghelli *et al.*, 1992, Matthes, 1992]. If this target system has orthogonal persistence, then the bulk values inherit that persistence. It has long been argued that during this mapping, the traditional database optimisations can be accomplished [Cooper, et al., 1987]. These optimisation techniques have yet to be convincingly demonstrated, though it still seems a valid proposition. A potential problem is whether the “standard” primitives are the appropriate target and whether their cost models are well enough developed to support optimisation.

The issue of whether to build-in or add-on bulk types is often posed [Matthes and Schmidt, 1991]. The resolution is unlikely to be simple. The addition of new bulk-types will always be required as it is inconceivable that all potential requirements could be anticipated. However, their definition in terms of primitives will remain a skilled craft. In an attempt to de-skill it to some extent and to provide a more appropriate target onto which all bulk-types could be mapped, a generic primitive has been suggested [Atkinson, et al., 1991].

Languages continue to be designed with built-in bulk types, e.g. Fibonacci. One reason for this is the issue of query notations. An alternative approach that is worth further exploration, is to support appropriate notations by using reflection to generate code from a suitable query notation.

Query languages and bulk types are intimately interrelated in the context of object models and persistent type systems. Queries can only be effectively expressed over collections of values or objects, which may be represented by bulk types. But aspects of the query, e.g. selection expressions, may require expressions based on any operation of any other type in the language, including other bulk types.

6.1.3 Transactions and Recovery

Typical persistent systems implement recovery well. Similarly typical database systems implement recovery effectively in the context of serialisable transactions. However, both show limitations when any more sophisticated form of transaction or concurrency is required.

Research is needed to validate the primitives and constructors that will allow a range of transaction types and concurrency models to be used with the same system [Krablin, 1987b, Stemple and Morrison, 1992, Munro, et al., 1994]. Again, generation and translation technology is anticipated to prevent the target becoming over-complicated and to relieve programmers from mechanisable detail. It remains a challenge to develop a recovery mechanism that has comprehensible behaviour in the context of the full range of transactions.

6.1.4 Evolutionary Constructs and Mechanisms

Infinite union or dynamic types, such as type **any**, provide the basic requirements to permit incremental changes to types. Reflection permits software to change its environment and to accommodate change in ways that are otherwise impossible without recourse to expensive interpretation of all operations. Exploration of how these may be combined to support system evolution has already made some progress.

The mechanisms for changing program parts are well developed in orthogonally persistent systems [Connor, 1990, Cutts, 1992, Sjøberg, 1993], as the executable code values can be transactionally replaced via store updates [Connor, et al., 1994b]. There are, however, many challenges that remain here.

- **To provide configuration and version managers:** the challenges here are, in part more complex than with traditional technology, since all types of data and associated types and programs need to be consistently versioned or configured. However, persistent technology may also come to our aid here. Its provision for all types of data and its support for reliable references, as well as its programming benefits combine to assist in the construction of these tools [Morrison *et al.*, 1994, Morrison, et al., 1995].
- **To provide convenient notations for specifying change:** One line of attack is to schema edit the higher-level data and process models and to generate the underlying changes. This runs into difficulties as programmers have often provided extra information during these mappings which it is infeasible to expect them to re-supply for each change. Another strategy is to support type editing on the types defined in the orthogonally persistent system itself. This pre-supposes that the mappings are comprehensible or it abandons generation from high-level notations.
- **To provide efficient propagation of change to the existing instances of types:** One strategy that works is to use a stable anchor for each type and accommodate extensions. This is very similar to the method commonly used for implementing object specialisations as fragments. However, it has costs of indirection and type projections. It can, in principle, be automated apart from generating new values and can be operated either in batch or incremental mode [Clamen, 1994]. But the trade-offs and implementation have yet to be explored in this context.
- **To assist programmers in managing change:** This requires recording or discovering dependencies, performing impact analysis and identifying the localities that need to be changed. Orthogonally persistent systems often hold the source code within the same regime as the types, data and executable code. It is therefore possible to search for these dependencies, and to present them as demonstrated by [Sjøberg, 1993]. The development of change management tools and methodologies for organising persistent object stores so that handling the consequences of change is tractable, is still in its infancy.
- **To provide appropriate change absorbers:** View mechanisms in databases act as change absorbers. That is, they are (usually manually) revised to prevent the propagation of change and the explicit use of a view also identifies very rapidly the cases where a program cannot possibly be affected by a schema change. The parsimonious use of explicit bindings between programs and stores, as typified by Napier88, also provides the rapid elimination of the unaffected programs, and abstract data types can, in principle, provide manually maintained change absorbers [Connor, et al., 1990]. Automated change absorption should be possible in many cases, and the relevant methods and mechanisms need development before they can be used for large scale applications [Atkinson, et al., 1993].

It would be unwise for any enterprise to seriously commit to a technology purporting to support persistence if they were not confident that it could adequately support schema change. Where persistent systems have imprecise type knowledge (e.g. those based on C and C++) it is doubtful whether there is enough information to safely support change.

6.1.5 Distribution and Autonomy

Persistent systems abstract over the locality of data, that is, the programmer does not need to be aware of whether it is on disk or in RAM. One approach to distributed orthogonally persistent systems is to continue this location transparency [Wai, 1989, Koch, et al., 1990, Gruber *et al.*, 1992, Daynès and Gruber, 1994, Dearle, et al., 1994]. This abstraction may eventually satisfy those that use distribution to achieve reliability and performance through distributing computation and migrating computation or data, but it depends on the discovery of adequate algorithms for

automating data and program movement. Although this approach offers desirable semantic simplicity it does not scale indefinitely, fails to accommodate change and does not have a realistic approach to failures [Mira da Silva *et al.*, 1995].

Others however wish to exploit their knowledge about processes and the network to write efficient distributed persistent systems. This requires that the locality of data be exposed [Liskov, 1988, Liskov *et al.*, 1990, Munro, 1993, Mira da Silva, 1995]. As yet, no tidy resolution of these conflicting requirements is in sight, and different researchers are pursuing the two models. *Obliq* is an interesting compromise [Cardelli, 1995] as it allows its programmers to indicate whether data or method should migrate but does not explicitly specify locations.

Generally, providers of PASs have to provide high levels of availability. However, distributed systems are prone to various forms of failure. It appears to be necessary to give programmers access to failure information in order that they may program this availability and inform users about the causes of failures and recovery expectations since the appropriate action is highly application dependent. But, to provide this information, the persistence abstraction has to be compromised. There is a challenge here: how to present to programmers a model of distribution which includes autonomy and failure but which does not leave them to handle all problems unaided.

An important reason for distribution is ownership. Different parts of the system belong to independent organisations who retain the right to make changes as they see fit. This autonomy is in fact independent of distribution and mechanisms that support it also contribute to scalability and the accommodation of change [Atkinson, *et al.*, 1993].

The crucial input that persistence brings to the more general search for adequate models of distribution and autonomy is a commitment to retaining simple, combinable primitives, even if this raises additional implementation challenges.

6.1.6 Scalable Systems

Over the lifetime of a PAS it is common to experience both massive growth and contraction, both of which may be difficult to predict. Often these variations in scale affect only parts of the PAS; for example, meta-data, data, program and users all accumulate around successful parts of a PAS while other parts may virtually atrophy. Technology is required that adapts well across all possible sizes of component collection. Although the range that can be efficiently accommodated is growing, and specific solutions have been found for very large bodies of data [Grossman, *et al.*, 1994], there is always a need for better adaptive algorithms.

However, there are limits to the performance that can be expected from algorithms that manage data on disk. For example, it is infeasible that disk garbage collection algorithms and store re-organisation algorithms will ever work well with arbitrarily large bodies of unconstrainably interconnected data. One suggested approach is to introduce logical partitions that assist with scalability from design through to store maintenance throughout the life-cycle. A complementary strategy is to utilise replication and parallelism in distributed systems.

Persistence researchers have a particularly difficult time verifying that their approaches to scale and evolution will actually work. It is very difficult to resource experiments of the required scale, realism and duration.

6.1.7 Integration with other Domains

Much of the safety and simplicity of orthogonally persistent systems is achieved by postulating a closed world, e.g. the well defined universe of discourse determined by the type system, or the well defined extent of meta-data, data and program determined by reachability. These totally closed worlds are unrealistic if transition is to be made to persistence. Indeed, at present, most persistent systems have connections with the external world (for example they can use UNIX files and UNIX shell commands).

More sophisticated interfaces are obviously needed. It has already been remarked that the work on mapping data models already supports some skill transfer. It could also be used to automatically generate bulk load and unload programs that are still type safe, for any of the data models. This could be extended to use the selective loading developed by Abiteboul and Cluet [Abiteboul *et al.*, 1993].

A data type complete RPC technology that uses the automatic generation of interfaces via reflection has been developed [Mira da Silva, 1995] but there are still problems to be overcome [Mira da Silva, *et al.*, 1995]. This kind of RPC can be extended to automatic translation at the interface and connection to processes in other languages. Indeed, as long ago as 1982 Buneman demonstrated how type-safe interfaces could be generated to external data collections [Buneman *et al.*, 1982]. More recently, Matthes has repeated that work in the context of a persistent system [Matthes, 1992]. Current persistent type systems permit these external data repositories to be safely modelled as abstract data types. Reflection allows these ADTs to be automatically generated when they are needed avoiding the high set up costs and the large name spaces of Buneman's mechanism.

Similarly, standard interfaces to a persistent system, or with an external system can, in principle, be generated. The extensive use of such generated, safe interfaces has not yet been investigated. It is not expected to raise conceptual difficulties beyond those raised by autonomy and distribution.

6.2 Exploiting Persistence

There are several ways in which orthogonally persistent system research may be exploited. Methodologies and tools can be developed that enhance the utility of the existing and envisaged systems. The concepts developed in one case can be re-applied elsewhere; in particular, programming languages and OODBMS can be modified to utilise the results.

6.2.1 Building and Observing Persistent Exemplars

The developers of orthogonally persistent systems believe that these systems are ready for serious evaluation. Wider use is needed to test their utility and to provide feedback to implementors so that the engineering issues can be better understood. Many of the issues are only apparent when the usage involves large volumes of program and data. It is crucial for effective evaluation and the generation of useful characterisations of load that the usage of these systems is realistic, in the sense that operational loads, data sets, programmers, users and changes are typical of real applications. Therefore, experiments should have sufficient resources to explore the critical issues at adequate scale and over sufficient elapsed time [Atkinson, 1992].

Good quality engineering depends on the systematic use of measurements. The persistent community therefore need to develop system models and measurement technology and practices so that PAS loads and the corresponding performance of PASs can be accurately described and compared, and so that design trade-offs can be based on reliable information [Atkinson *et al.*, 1992, Scheuerl, *et al.*, 1995].

6.2.2 Using Type-safe Linguistic Reflection

Reflection is especially significant in persistent systems as the program generation costs can be amortised over more applications using persistent memoisation and because it is the basis for system evolution. It has two important roles: it allows adaptive behaviour to be programmed and it facilitates the writing of processors from specialised notations to a target persistent language. Both techniques have already been well demonstrated. Two lines of research are identifiable.

- Research to improve the notations and libraries available for those writing reflexive programs is already underway [Cooper & Kirby, 1994, Kirby, et al., 1994]. Further elimination of the noise generated by embedding representations of a program within programs is probably relatively easily achieved. The more challenging task is to help the programmer think simultaneously and accurately about two computations: the computation to form the new code and the computation that code will eventually perform. Part of that challenge may be to make more apparent the way the environments of the two computations relate.
- Research is also needed to explore the full potential of reflection. Examples that have been mentioned already are the construction of optimised strategies for computations over bulk values based on observed properties of those values and the simplification of target persistent languages given the ready availability of reflection. The research into the provision of data models and user interface generators is well underway. Research into the use of reflection for generating safe external interfaces will continue. The use of reflection to automate the management of the consequences of change is a new area ripe for exploration. These are just a few of the possible applications of reflection deserving exploration.

Given the advent of type-safe reflection and tools that assist programmers to use it, there is likely to be an explosion of applications. While these may be important in themselves, it is perhaps even more important to extract and identify effective methodologies for reflection in order to enable other programmers to use it well.

6.2.3 Methodologies and Tools

Constructing application systems using orthogonal persistence allows new techniques and new structures. For example, as structural information is neither lost nor obscured if it is left in the store by one program to be used by another, there is a trend towards partitioning applications into smaller units. As another example, libraries and application programs can both be built incrementally in the persistent store using the same transactional mechanisms as the application uses. This permits incremental construction and replacement of small units of program and data, even while the system is in use.

As experience is gained with writing PASs, methodologies appropriate for this new technology are emerging [Dearle, 1989, Connor, 1990, Sjøberg, 1993]. These require further development and need to be made accessible to application programmers in a variety of persistent technologies including OODBMS through the provision of tutorials, tools and exemplars.

Crucial to making these methodologies usable is the provision of an appropriate persistent programming environment. Three experiments are underway in this direction at present: the hyper-programming workbench [Connor *et al.*, 1994a], the PIPE programming environment [Dearle *et al.*, 1992a] and the persistent workshop [Sjøberg, et al., 1995]. The hyper-programming workbench is allowing experiments with new construction techniques and with construction-time type-checking. The persistent workshop already provides a set of tools to help programmers, including:

- specialised editors, automated binding resolution and compilers;
- library management tools, including aids to finding components based on information retrieval techniques; and
- visualisers that are intended to help programmers understand the state of the persistent store.

These tools will continue to develop as programming aids. The tools which derive dependencies and automate re-compilations and re-executions that are under development lie at the boundary between supporting individual programming and supporting software engineering. If a methodology has been accurately followed,

these tools can be more precise. Consequently, the facilities for verifying consistency with a methodology are under development [Sjøberg *et al.*, 1994a]. The methodology has been chosen in the belief that it will facilitate the change processes that dominate a PAS life-cycle. This has yet to be verified, and even the advice to the programmers on the consequences of change is still in its infancy [Sjøberg *et al.*, 1993, Sjøberg *et al.*, 1994b].

6.2.4 Influencing other Technologies

The results achieved by applying the persistent design principles can be applied in other domains. Three such areas where persistence results could have a beneficial influence are: programming languages, OODBMS and operating systems.

Existing commercial programming languages (C, C++, Ada, etc.) are commonly used in building PASs. There is a natural desire to produce persistent versions. The persistence results show that this can only be done safely if the types of each value are known unambiguously. It is therefore probable that it will be worthwhile establishing persistent dialects of these languages *that have this safety property*.

From time-to-time new programming languages emerge. The designers of those languages will have omitted important and feasible facilities if they fail to include orthogonal persistence.

Object-oriented databases are following a similar path to that followed by persistent technology. The persistent experience can guide OODBMS design and implementation. For example, basing persistence on reachability from identified persistent roots and providing orthogonal persistence, now appears obvious and has been adopted in O_2 [Bancilhon *et al.*, 1992]. OODBMS would also benefit from properly incorporating the stored program (methods) into the persistent store, so that consistent bindings between code and data can be maintained [Morrison, et al., 1995].

Operating systems already provide non-orthogonal (weakly-typed) persistence in the form of file stores. It is apparent from the complexities of using their environments (e.g. **setenv** in UNIX) that orthogonal persistence and better defined models of binding to the persistent store would be a significant advantage. Some work is also underway on persistent operating systems with a view to having such facilities and supporting persistent programming generally in a much more efficient way. Certainly, the requirements of orthogonal persistence place new demands on the low-level mechanisms provided by operating systems. Persistence research should therefore influence the way operating systems present their storage facilities and the facilities they offer to system implementors. But persistence may also have another role to play, as the harmonisation between databases and programming languages that has been its main focus to date, could also be applied to the relationship between programs and the operating systems they use. Why not have the same naming, binding and typing rules when using operating system facilities (files, system calls) as are used in the language? Of course, as was the case with databases, both would ultimately evolve to meet their common requirements.

6.3 Delivering Persistence

Although reasonable scale and performance can be obtained from current orthogonally persistent technology, there is a continuing search for better engineering. The challenge is to develop persistent object stores that provide the full range of facilities for collections of data from a few thousand bytes to terabytes.

Four models of providing persistence have been presented.

- **Provision of a “library” of persistent facilities in a standard language:**
This may be a useful pragmatic intermediate step but we believe it is unlikely to lead to safe and comprehensible PAS building facilities.

- **Combination of a standard language and a separate store:** This is commercially a most active line of attack at present as various OODBMS vendors combine their product with various languages. In reality they often modify the implementation of the language and adapt the provisions of their store. This is therefore a valid intermediary process. It may be difficult for it to give satisfactory results in the long term due to the lack of reliable type information and inconsistent behaviour when the system is stressed. Certainly, the potential of this approach to support PAS evolution should be investigated *before* it becomes commercially critical.
- **Integrated design and provision:** This approach has been taken by the persistent programming community. It is clearly a radical and therefore risky departure from commercial practice. However, this technology is at least as mature as relational technology was when it was first explored commercially.
- **Re-building the whole support system:** This approach starts with the construction of a persistent operating system (or even hardware) and is therefore a much longer term programme of research. Without its results, persistent support systems will continue to conflict with operating systems at a cost of performance and possibly of reliability. This avenue of research is therefore important as it will develop technology which will be required as persistent systems become prevalent.

Only by building and using operational quality systems taking one of the last two forms above can the real potential of orthogonally persistent systems be verified [Atkinson, 1992].

6.3.1 Efficient Object Stores

Three strategies are in wide use to provide persistent object stores:

- Constructing above standard operating systems using block transfers between files. This has the advantage of portability.
- Constructing above standard operating systems using memory mapping technology. This exploits paging hardware but has less locality and may suffer from slow operating system interfaces and inappropriate protection granularity. It can have the advantage of running programs written in a non-persistent language in a persistent mode [Singhal, et al., 1992].
- Constructing above light-weight or specialised operating systems using memory mapping. In some cases these experiments also exploit specialised hardware. Very little portability is available and performance gains are as yet unproven.

Orthogonal to the above class of design decisions is the extent to which the long-term form of an object differs from its active-computational form. For example, pointers can be radically different [Koch, et al., 1990, Moss, 1990, Suzuki *et al.*, 1994, Kemper and Kossmann, 1995], code can be target independent [Atkinson, et al., 1993] and data may be compressed and encrypted.

There are many other variations in detail that are currently being explored, for example, the mechanisms for recovery [Scheuerl, et al., 1995]. It is clear that the space of alternative store designs is very large and is, as yet, poorly charted. Research is needed to develop reliable information to guide store implementors and to provide high-performance stores with predictable behaviour over a range of loads.

The situation is much more complex when forms of distributed persistent store are attempted. There are some experiments in client-server technology [Dearle *et al.*, 1992b, Dearle, et al., 1994] but there is, as far as the authors know, little work yet on generally distributed stores on a cluster of well-connected machines.

6.3.2 Efficient Bulk Types

The requirements for bulk types is well established [Atkinson and Buneman, 1987] and their implementation has received much attention from the database community. Their key achievements should be transferable to the persistent context. However, this may result in pressure to change aspects of the store implementation and even language primitives. Some compromise between totally add-on and totally built-in provision is likely to prevail eventually. Full scale optimisation cannot be implemented without better mathematical models of store behaviour. These need to be vigorously sought and experimentally validated.

6.3.3 Code Generation for Reflexive Higher-Order Persistent Languages

There are several difficulties with code generation for higher-order languages. For example, it is convenient to use some intermediate target language such as C to avoid computer architecture specificity. This requires techniques for handling higher-order procedures such as those in [Appel, 1992]. But orthogonal persistence requires that the generated code reside in the transactional stable store. Consequently code (or more precisely, procedure closures) must be represented as objects that can be shipped in and out of store, garbage collected and stabilised [Atkinson and Morrison, 1985]. Furthermore, these procedures may be generated during execution as a consequence of run-time reflection. Once distribution or longer-term persistence is required, these procedures must be loadable onto any architecture. This combination of requirements is challenging. Consequently, only prototype code generators exist at present that meet all these requirements [Bushell *et al.*, 1994].

6.3.4 Pre-populated Stores

A persistent store can be shipped in pre-populated form, just as many object-oriented languages and OODBMS are (a good example is the SmallTalk persistent virtual image). The challenge is to determine what should go into this initial population and also to find ways of organising it and accessing it that enable the majority of programmers to make good use of its facilities. Not only will this pre-population contain an extensive library of useful code but it can also contain useful data structures (e.g. tables of unit conversion constants, cartographic images, pictures, etc.). Furthermore, the range of code that can be usefully shipped includes access algorithms (B⁺-trees, R^{*}-trees, etc.) which can be already tuned to match the properties of the shipped store. The development and consideration of these issues can be conducted by a wide community of researchers and eventually we should have the PAS equivalent of good numerical algorithms libraries with much more convenient mechanisms for using them.

Typically new releases of these store populations are built, and users who are already using an earlier version may want to preserve their work but also avail themselves of the new libraries. This is just one example of a collection of problems concerned with merging, replacing and shipping subsets of the objects in one store to other stores.

7 Summary and Conclusions

An overview of orthogonally persistent systems has been presented starting from their motivation and design principles. These achieve simplicity from consistency and regularity. The lack of exceptions enhances the power while the reduction of issues to be dealt with by programmers reduces application design, construction and maintenance costs.

Languages and systems that achieve orthogonal persistence have been reported. Crucial to their achievement of integrating database and programming requirements is the recognition of the close relationship between schema and types. The challenge of obtaining the correct balance between precision and description is addressed by the introduction of type operators and the concomitant use of polymorphism. This allows adaptable type systems to be developed with power similar to existing data models

and an extensibility necessary for long-term service [Connor *et al.*, 1991, Connor & Morrison, 1992].

The introduction of dynamic checking points within the type system serves three purposes: it allows type checking to be incremental and therefore feasible in large applications; it allows incremental binding; and it permits localisation of schema changes to subgraphs delimited by these points. The support for incremental evolution of persistent applications requires new binding mechanisms that utilise the dynamic checking and provide consistent naming schemes throughout applications.

The development of constructs that provide concurrency, transactions and recovery taken together still lacks unifying concepts that cover the full range of possibilities from those present and sought in databases to those available in programming languages. However, solutions exist for important subdivisions of that space and the present lines of research are rapidly expanding these.

Much research has focused on how to support orthogonally persistent systems. Four architectures are all being explored. The least committing involves extending an existing system with libraries and possibly pre-processors. Although this can provide short term benefits with few risks, it has fundamental limitations in the longer term. A common strategy is to combine two existing systems. This has operated since the early days of databases but the degree of integration has advanced considerably with the advent of OODBMS. It may still leave programmers coping with two models, particularly when they are concerned with failure semantics and other symptoms of system stress. They may also find that they still have to perform mappings between representations, though this is now much diminished.

Recently, integrated systems have been constructed in standard computing environments. These achieve the desired consistency and are portable. They are at the very least a necessary intermediary on the path to wider use of orthogonally persistent systems and are likely to have long-term utility. Their only drawback is that they operate in an environment tuned to a different style of computation and therefore may have efficiency problems due to this mismatch. Although this may present challenges to current implementors it is feasible that the supporting operating systems and hardware can evolve to meet the new needs of this architecture.

The final architecture may be viewed as prototyping this evolution or as a platform for a more radical transition. In this architecture the underlying operating systems and even the hardware are re-designed and re-implemented to provide orthogonal persistence for all the systems they support. This radical replacement of the persistent support system would, if it were eventually adopted, have the benefit of ensuring consistent persistent behaviour, even under stress, of all the software and data it supported simplifying the tasks of programmers and users.

In all of the last three avenues of research much effort goes into object store design. The challenge is to find mechanisms for object movement, representation, translation, concurrency, transactions, recovery and space management that combine well together on the available hardware and software platforms. The size and load on parts of a persistent system is likely to vary dramatically between applications and during the lifetimes of some applications. Consequently, there is considerable interest in adaptive algorithms. Persistent object store technology is now quite sophisticated and can deliver reasonable performance and full functionality over specific areas of this design space.

Type-safe linguistic reflection has been developed extensively in the context of persistent systems for three reasons: it enables systems to evolve; it allows specialised notations to be used from which equivalent types and programs are generated in the persistent language; and it permits safe computations over types yielding greater genericity.

Using these technologies, substantial applications have been built. Their construction and maintenance uses incremental techniques that originated with database

programming but finer grain and more general incrementality is now possible. Among the applications are tools and programming environments that support persistent programming and exploit the potential of persistence to give reliable access to contextual information. Hyper-programming is supported in such an environment. This is a new technique for program construction that forms static bindings from the source program representation to values in the store while the source is being edited.

The present uses of persistence for constructing programming environments and software engineering tools not only demonstrate the viability of persistent technology but also its ability to support the introduction of new tools, flexibly coupled via a store that preserves all structural information and affords continuous type safety. The software engineering tools particularly exploit the reliable references and the data modelling tools make substantial use of reflection.

Many avenues of research are being pursued in persistent systems and many more are opening up. The three main avenues of research are: the extension of persistence results to other contexts; the exploitation of the potential of persistence to facilitate new applications with sophisticated computational requirements; and the search for improved engineering for persistent support systems.

After fifteen years of research, orthogonally persistent systems have reached the point where they should be used by a wider community. Researchers who build applications with complex long-term data may accelerate their experiments by using one of the more mature persistent systems. Programming language, operating system and database system designers should at least examine the results obtained with orthogonal persistence with a view to incorporating the best and most relevant of them into their own designs. Ideally, one of the mature research systems will be commercially supported so that this new technology can be seriously evaluated under realistic use.

8 Acknowledgements

The authors acknowledge the support of the European Union through FIDE₂ (ESPRIT Basic Research Action 6309) and of the British EPSRC through grants: EPSRC Grant GR/J67611 and DRASTIC EPSRC Grant GR/J99285. The authors are very grateful for the generous help in preparing the paper given by Graham Kirby, Miguel Mira da Silva, Quintin Cutts, Paul Philbrow and Ray Welland. The following also made helpful comments on this paper: Peter Bailey, Richard Connor, Richard Cooper, Peter Dickman, Huw Evans, Stewart Macneill, Dave Munro, Susan Spence and Cathy Waite.

9 Bibliography

- Abiteboul, S., Cluet, S. & Milo, T., 1993. Querying and Updating the File. *In Proc. 19th International Conference on Very Large Data Bases*, Dublin, Ireland pp 73-84.
- Agrawal, R. & De Witt, D., 1985. Recovery Architectures for Multiprocessor Database Machines. *In Proc. SIGMOD International Conference on Management of Data* pp 131-147.
- Akima, N. & Ooi, F., 1989. Industrializing Software Development: A Japanese Approach. *IEEE Software* 6, 2 pp 13-21.
- Albano, A., Bergamini, R., Ghelli, G. & Orsini, R., 1993. An Object Data Model with Roles. *In Proc. 19th International Conference on Very Large Data Bases*, Dublin, Ireland pp 39-51.
- Albano, A., Bergamini, R., Ghelli, G. & Orsini, R., 1995. An Introduction to the Database Programming Language Fibonacci. *Journal of Very Large Databases* 4, 3.
- Albano, A., Cardelli, L. & Orsini, R., 1985. Galileo: a Strongly Typed, Interactive Conceptual Language. *ACM Transactions on Database Systems* 10, 2 pp 230-260.
- Albano, A., Dearle, A., Ghelli, G., Marlin, C.D., Morrison, R., Orsini, R. & Stemple, D., 1989. A Framework for Comparing Type Systems for Database Programming Languages. *In Database Programming Languages*, Hull, R., Morrison, R. & Stemple, D. (ed.), Morgan Kaufmann pp 170-178.
- Albano, A. & Morrison, R. (ed.), 1992. *Persistent Object Systems: Implementation and Use*. Springer-Verlag.
- Appel, A.W., 1992. *Compiling with Continuations*. Cambridge University Press, Cambridge, UK.
- Atkinson, M.P., 1978. Programming Languages and Databases. *In Proc. 4th IEEE International Conference on Very Large Databases* pp 408-419.
- Atkinson, M.P., 1992. Persistent Foundations for Scalable Multi-Paradigm Systems. *In Proc. International Workshop on Distributed Object Management*, Edmonton, Canada pp 26-50.
- Atkinson, M.P., Bailey, P., Christie, D., Cropper, K. & Philbrow, P., 1993. Towards Bulk Type Libraries for Napier88 . ESPRIT BRA Project 6309 FIDE₂ Technical Report FIDE/93/78.
- Atkinson, M.P., Bailey, P.J., Chisholm, K.J., Cockshott, W.P. & Morrison, R., 1983. An Approach to Persistent Programming. *Computer Journal* 26, 4 pp 360-365.
- Atkinson, M.P., Bailey, P.J., Chisholm, K.J., Cockshott, W.P. & Morrison, R., 1984. Progress with Persistent Programming. *In Database, Role and Structure*, Stocker, P.M., Atkinson, M.P. & Gray, P.M. (ed.), Cambridge University Press pp 245-310.

- Atkinson, M.P., Birnie, A., Jackson, N. & Philbrow, P.C., 1992. Measuring Persistent Object Systems. In *Persistent Object Systems*, Albano, A. & Morrison, R. (ed.), Springer-Verlag, Proc. 5th International Workshop on Persistent Object Systems, San Miniato, Italy pp 63-85.
- Atkinson, M.P. & Buneman, O.P., 1987. Types and Persistence in Database Programming Languages. *ACM Computing Surveys* 19, 2 pp 105-190.
- Atkinson, M.P., Buneman, O.P. & Morrison, R., 1985. Proceedings of the Persistence and Data Types Workshop, Appin 85 . Universities of Glasgow and St Andrews Technical Report PPRR-16-85.
- Atkinson, M.P., Buneman, O.P. & Morrison, R., 1987. Persistent Object Systems: their design, implementation and use, Appin 87 . Universities of Glasgow and St Andrews Technical Report PPRR-44-87.
- Atkinson, M.P., Buneman, O.P. & Morrison, R., 1988a. Binding and Type Checking in Database Programming Languages. *Computer Journal* 31, 2 pp 99-109.
- Atkinson, M.P., Buneman, O.P. & Morrison, R. (ed.), 1988b. *Data Types and Persistence*. Springer-Verlag, New York.
- Atkinson, M.P., Chisholm, K.J. & Cockshott, W.P., 1982. PS-algol: An Algol with a Persistent Heap. *ACM SIGPLAN Notices* 17, 7 pp 24-31.
- Atkinson, M.P., Chisholm, K.J. & Cockshott, W.P., 1983. CMS - A Chunk Management System. *Software – Practice and Experience* 13, 3 pp 259-272.
- Atkinson, M.P., Lécluse, C., Philbrow, P. & Richard, P., 1991. Design Issues in a Map Language. In *Bulk Types & Persistent Data*, Kanellakis, P. & Schmidt, J.W. (ed.), Morgan Kaufmann pp 20-32.
- Atkinson, M.P., Maier, D. & Benzaken, V. (ed.), 1995. *Persistent Object Systems, Tarascon 1994*. Springer-Verlag.
- Atkinson, M.P. & Morrison, R., 1985. Procedures as Persistent Data Objects. *ACM Transactions on Programming Languages and Systems* 7, 4 pp 539-559.
- Atkinson, M.P. & Morrison, R., 1986. Integrated Persistent Programming Systems. In *Proc. 19th International Conference on Systems Sciences*, Hawaii pp 842-854.
- Atkinson, M.P. & Morrison, R., 1987. Polymorphic Names, Types, Constancy and Magic in a Type Secure Persistent Object Store. In *Proc. 2nd International Workshop on Persistent Object Systems*, Appin, Scotland pp 1-12.
- Atkinson, M.P. & Morrison, R., 1989. Coordinators of Persistent Systems Track. In *Proc. 22nd International Conference on Systems Sciences*, Hawaii pp 17-18.
- Atkinson, M.P., Morrison, R. & Pratten, G.D., 1986. Designing a Persistent Information Space Architecture. In *Proc. 10th IFIP World Congress*, Dublin pp 115-120.
- Atkinson, M.P., Sjøberg, D.I.K. & Morrison, R., 1993. Managing Change in Persistent Object Systems. In *Proc. JSSST International Symposium on Object Technologies for Advanced Software*, Kanazawa, Japan pp 315-338.

- Bancilhon, F., Barbedette, G., Benzaken, V., Delobel, C., Gamerman, S., Lécluse, C., Pfeffer, P., Richard, P. & Valez, F., 1988. The Design and Implementation of O_2 , an Object-Oriented Database System. In *Lecture Notes in Computer Science 334*, Dittrich, K.R. (ed.), Springer-Verlag pp 1-22.
- Bancilhon, F. & Buneman, O.P. (ed.), 1990. *Advances in Database Programming Languages*. Addison-Wesley and ACM Press.
- Bancilhon, F., Delobel, C. & Kanellakis, P. (ed.), 1992. *The Story of O_2 : Building an Object-Oriented Database System*. Morgan Kaufmann.
- Benzaken, V., Delobel, C. & Harrus, G., 1991. Clustering Strategies in the O_2 Object-Oriented Database System . ESPRIT BRA Project 3070 FIDE Technical Report FIDE/91/21.
- Berman, S., 1991. P-Pascal: A Data-Oriented Persistent Programming Language. PhD Thesis, University of University of Cape Town.
- Bocca, J. & Bailey, P.J., 1987. Logic Languages and Relational DBMSs—the Point of Convergence. In *Proc. 2nd International Workshop on Persistent Object Systems*, Appin, Scotland pp 346-362.
- Bott, M.F. (ed.), 1989. *ECLIPSE: An Integrated Project Support Environment*. Peter Peregrinus.
- Breazu-Tannen, V., Buneman, O.P. & Naqvi, S., 1991. Structural Recursion as a Query Language. In *3rd International Workshop on Database Programming Languages*, Kanellakis, P. & Schmidt, J.W. (ed.), Morgan Kaufmann, Proc. 3rd International Workshop on Database Programming Languages, Nafplion, Greece pp 9-19.
- Bretl, B., Maier, D., Otis, A., Penney, J., Schuchardt, B., Stein, J., Williams, E.H. & Williams, M., 1989. The GemStone Data Management System. In *Object-Oriented Concepts, Databases and Applications*, Kim, W. & Lochovsky, F. (ed.), ACM Press and Addison Wesley pp 283-308.
- Brookes, S.D., Hoare, C. & Roscoe, A., 1980. A Theory of Communicating Sequential Processes . Carnegie-Mellon University Technical Report CMU-CS-83-153.
- Brown, A.L., 1987. A Distributed Stable Store . Universities of Glasgow and St Andrews Technical Report PPRR-50-87.
- Brown, A.L., 1989. Persistent Object Stores. Ph.D. Thesis, University of St Andrews.
- Brown, A.L. & Cockshott, W.P., 1985. The CPOMS Persistent Object Management System . Universities of Glasgow and St Andrews Technical Report PPRR-13-85.
- Brown, A.L., Mainetto, G., Matthes, F., Müller, R. & McNally, D.J., 1992. An Open System Architecture for a Persistent Object Store. In *Proc. 25th International Conference on Systems Sciences*, Hawaii pp 766-776.
- Brown, A.L. & Morrison, R., 1992. A Generic Persistent Object Store. *Software Engineering Journal* 7, 2 pp 161-168.

- Brown, J.C., 1993. A Library Explorer for the Napier88 Glasgow Libraries. MSc Thesis, University of University of Glasgow.
- Bruynooghe, R.F., Parker, J.M. & Rowles, J.S., 1991. PSS: A System for Process Enactment. *In Proc. 1st International Conference on the Software Process: Manufacturing Complex Systems.*
- Buneman, O.P., Hirschberg, J. & Root, D., 1982. Integrating CODASYL with High Level Programming Languages. *In Proc. 2nd British National Conference on Databases*, Bristol, England.
- Buneman, O.P., Libkin, L., Suciu, D., Tannen, V. & Wong, L., 1994. Comprehension Syntax. *ACM SIGMOD Record* 23, 1 pp 87-96.
- Burstall, R., 1984. Programming with Modules as Typed Functional Programming. *In Proc. International Conference on 5th Generation Computing Systems.*
- Burstall, R.M., Collins, J.S. & Popplestone, R.J., 1971. *Programming in POP-2*. Edinburgh University Press, Edinburgh, Scotland.
- Bushell, S.J., Dearle, A., Brown, A.L. & Vaughan, F.A., 1994. Using C as a Compiler Target Language for Native Code Generation in Persistent Systems. *In Persistent Object Systems, Tarascon 1994*, Atkinson, M.P., Maier, D. & Benzaken, V. (ed.), Springer-Verlag, Proc. 6th International Workshop on Persistent Object Systems, Tarascon, France pp 164-183.
- Campbell, R.H. & Haberman, A.N., 1974. The Specification of Process Synchronisation by Path Expressions. *In Lecture Notes in Computer Science 16*, Springer-Verlag.
- Cardelli, L., 1986. Amber. *In Lecture Notes in Computer Science 242*, Springer-Verlag pp 21-47.
- Cardelli, L., 1989. Typeful Programming . DEC Technical Report 45.
- Cardelli, L., 1995. Obliq: A Language with Distributed Scope. *Computing Systems* 8, 1 pp 27-59.
- Cardelli, L. & MacQueen, D.B., 1988. Persistence and Type Abstraction. *In Data Types and Persistence*, Atkinson, M.P., Buneman, O.P. & Morrison, R. (ed.), Springer-Verlag pp 31-42.
- Cardelli, L. & Wegner, P., 1985. On Understanding Types, Data Abstraction and Polymorphism. *ACM Computing Surveys* 17, 4 pp 471-523.
- Chan, A., Dayal, U. & Fox, S., 1987. An Ada-Compatible Distributed Database Management System. *In Proceedings of the IEEE, Special Issue on Distributed Databases*, IEEE pp 674-694.
- Chiu, S.-Y. & Levin, R., 1993. The Vesta Repository: A File System Extension for Software Development . DEC Systems Research Center Technical Report 106.
- Clamen, S.M., 1994. Schema Evolution and Integration. *Distributed and Parallel Databases* 2, 1 pp 101-126.
- Cluet, S. & Delobel, C., 1991. Towards a Unification of Rewrite Based Optimisation Techniques for Object-Oriented Queries . ESPRIT BRA Project 3070 FIDE Technical Report FIDE/91/19.

- Cluet, S. & Moerkotte, G., 1993. Nested Queries in Object Bases . ESPRIT BRA Project 6309 FIDE₂ Technical Report FIDE/93/69.
- Cockshott, W.P., Atkinson, M.P., Chisholm, K.J., Bailey, P.J. & Morrison, R., 1984. POMS: A Persistent Object Management System. *Software – Practice and Experience* 14, 1 pp 49-71.
- Connor, R.C.H., 1990. Types and Polymorphism in Persistent Programming Systems. Ph.D. Thesis, University of St Andrews.
- Connor, R.C.H., Atkinson, M.P., Berman, S., Cutts, Q.I., Kirby, G.N.C. & Morrison, R., 1993. The Joy of Sets. In *Database Programming Languages*, Beerl, C., Ohori, A. & Shasha, D.E. (ed.), Springer-Verlag, Proc. 4th International Conference on Database Programming Languages, New York City pp 417-433.
- Connor, R.C.H., Brown, A.B., Cutts, Q.I., Dearle, A., Morrison, R. & Rosenberg, J., 1990. Type Equivalence Checking in Persistent Object Systems. In *Implementing Persistent Object Bases, Principles and Practice*, Dearle, A., Shaw, G.M. & Zdonik, S.B. (ed.), Morgan Kaufmann, Proc. 4th International Workshop on Persistent Object Systems, Martha's Vineyard, USA pp 151-164.
- Connor, R.C.H., Cutts, Q.I., Kirby, G.N.C., Moore, V.S. & Morrison, R., 1994a. Unifying Interaction with Persistent Data and Program. In *Interfaces to Database Systems*, Sawyer, P. (ed.), Springer-Verlag, Proc. 2nd International Workshop on User Interfaces to Databases, Ambleside, Cumbria, 1994 pp 197-212.
- Connor, R.C.H., Cutts, Q.I., Kirby, G.N.C. & Morrison, R., 1994b. Using Persistence Technology to Control Schema Evolution. In *Proc. 9th ACM Symposium on Applied Computing*, Phoenix, Arizona pp 441-446.
- Connor, R.C.H., Dearle, A., Morrison, R. & Brown, A.L., 1990. Existentially Quantified Types as a Database Viewing Mechanism. In *Lecture Notes in Computer Science 416*, Bancilhon, F., Thanos, C. & Tsichritzis, D. (ed.), Springer-Verlag, Proc. 2nd International Conference on Extending Database Technology, Venice, Italy pp 301-315.
- Connor, R.C.H., McNally, D.J. & Morrison, R., 1991. Subtyping and Assignment in Database Programming Languages. In *3rd International Workshop on Database Programming Languages*, Kanelakis, P. & Schmidt, J.W. (ed.), Morgan Kaufmann, Proc. 3rd International Workshop on Database Programming Languages, Nafplion, Greece.
- Connor, R.C.H. & Morrison, R., 1992. Subtyping Without Tears. In *Proc. 15th Australian Computer Science Conference*, Hobart, Tasmania pp 209-225.
- Cook, J.E., Wolf, A.L. & Zorn, B.G., 1993. The Performance of Partitioned Garbage Collection in Object Databases . University of Colorado Technical Report CU-CS-653-93.
- Cooper, R. & Kirby, G.N.C., 1994. Type-Safe Linguistic Run-time Reflection - A Practical Perspective. In *Persistent Object Systems, Tarascon 1994*, Atkinson, M.P., Maier, D. & Benzaken, V. (ed.), Springer-Verlag, Proc. 6th International Workshop on Persistent Object Systems, Tarascon, France pp 332-354.

- Cooper, R.L., 1990a. Configurable Data Modelling Systems. *In Proc. 9th International Conference on the Entity Relationship Approach*, Lausanne, Switzerland pp 35-52.
- Cooper, R.L., 1990b. On The Utilisation of Persistent Programming Environments. Ph.D. Thesis, University of Glasgow.
- Cooper, R.L., Atkinson, M.P., Dearle, A. & Abderrahmane, D., 1987. Constructing Database Systems in a Persistent Environment. *In Proc. 13th International Conference on Very Large Data Bases* pp 117-125.
- Cooper, R.L. & Qin, Z., 1992. A Graphical Data Modelling Program With Constraint Specification and Management. *In Proc. 10th British National Conference on Databases*, Aberdeen pp 192-208.
- Cooper, R.L. & Qin, Z., 1994. A Generic Data Model for the Support of Multiple User Interaction Facilities. In *Lecture Notes in Computer Science 881*, Loucopoulos, P. (ed.), Proc. International Conference on the Entity Relationship Approach pp 351-368.
- Currie, I.F., 1985. Filestore and Modes in Flex. *In Proc. 1st International Workshop on Persistent Object Systems*, Appin, Scotland pp 325-334.
- Curtis, B., Kellner, M.I. & Over, J., 1992. Process Modeling. *Communications of the ACM* 35, 9 pp 75-90.
- Cutts, Q.I., 1992. Delivering the Benefits of Persistence to System Construction and Execution. Ph.D. Thesis, University of St Andrews.
- Cutts, Q.I., Connor, R.C.H., Kirby, G.N.C. & Morrison, R., 1994. An Execution Driven Approach to Code Optimisation. *In Proc. 17th Australasian Computer Science Conference*, Christchurch, New Zealand pp 83-92.
- Dahl, O. & Nygaard, K., 1966. Simula, an Algol-Based Simulation Language. *Communications of the ACM* 9, 9 pp 671-678.
- Davie, A.J.T. & McNally, D.J., 1990. Statically Typed Applicative Persistent Language Environment (STAPLE) Reference Manual . University of St Andrews Technical Report CS/90/14.
- Daynès, L. & Gruber, O., 1994. Customizing Concurrency Controls using Graph of Locking Capabilities. In *Persistent Object Systems, Tarascon 1994*, Atkinson, M.P., Maier, D. & Benzaken, V. (ed.), Springer-Verlag, Proc. 6th International Workshop on Persistent Object Systems, Tarascon, France pp 147-161.
- Dearle, A., 1988. On the Construction of Persistent Programming Environments. Ph.D. Thesis, University of St Andrews.
- Dearle, A., 1989. Environments: A flexible binding mechanism to support system evolution. *In Proc. 22nd International Conference on Systems Sciences*, Hawaii pp 46-55.
- Dearle, A. & Brown, A.L., 1988. Safe Browsing in a Strongly Typed Persistent Environment. *Computer Journal* 31, 6 pp 540-544.

- Dearle, A., Cutts, Q.I. & Kirby, G.N.C., 1990. Browsing, Grazing and Nibbling Persistent Data Structures. In *Persistent Object Systems*, Rosenberg, J. & Koch, D.M. (ed.), Springer-Verlag, Proc. 3rd International Workshop on Persistent Object Systems, Newcastle, Australia pp 56-69.
- Dearle, A., di Bona, R., Farrow, J., Henskens, F., Lindström, A., Rosenberg, J. & Vaughan, F., 1994. Grasshopper: An Orthogonally Persistent Operating System. *Computer Systems* 7, 3 pp 289-312.
- Dearle, A., Marlin, C.D. & Dart, P., 1992a. A Hyperlinked Persistent Software Development Environment. In *Proc. Hyper-Oz '92: A Workshop on Hypertext Activities in Australia*, Adelaide, Australia.
- Dearle, A., Rosenberg, J., Henskens, F.A., Vaughan, F. & Maciunas, K.J., 1992b. An Examination of Operating System Support for Persistent Object Systems. In *Proc. 25th International Conference on System Sciences*, Hawaii pp 779-789.
- Dearle, A., Shaw, G.M. & Zdonik, S.B. (ed.), 1990. *Implementing Persistent Object Bases: Principles and Practice*. Morgan Kaufmann.
- Deux, O., 1990. The Story of O₂. *IEEE Transactions on Knowledge and Data Engineering* 2, 1.
- Deux, O., 1991. The O₂ System. *Communications of the ACM* 34, 10 pp 34-48.
- Dijkstra, E.W., 1968a. Cooperating Sequential Processes. In *Programming Languages*, Genuys, F. (ed.), Academic Press pp 43-112.
- Dijkstra, E.W., 1968b. The Structure of the T.H.E. Multiprogramming System. *Communications of the ACM* 11, 5 pp 341-346.
- Ellis, C.A. & Gibbs, S.J., 1990. Concurrency Control in Groupware Systems. In *Proc. ACM-SIGMOD International Conference on Management of Data* pp 399-407.
- Eswaran, K.P., Gray, J.N., Lorie, R.A. & Traiger, I.L., 1976. The Notions of Consistency and Predicate Locks in a Database System. *Communications of the ACM* 19, 11 pp 624-633.
- Farkas, A. & Dearle, A., 1993. Octopus: A Reflective Language Mechanism for Object Manipulation. In *Database Programming Languages*, Beerli, C., Ohori, A. & Shasha, D.E. (ed.), Springer-Verlag, Proc. 4th International Conference on Database Programming Languages, New York City pp 50-64.
- Farkas, A.M., Dearle, A., Kirby, G.N.C., Cutts, Q.I., Morrison, R. & Connor, R.C.H., 1992. Persistent Program Construction through Browsing and User Gesture with some Typing. In *Persistent Object Systems*, Albano, A. & Morrison, R. (ed.), Springer-Verlag, Proc. 5th International Workshop on Persistent Object Systems, San Miniato, Italy pp 376-393.
- Fegaras, L., Sheard, T. & Stemple, D., 1992. Uniform Traversal Combinators: Definition, Use and Properties. In *Proc. 11th International Conference on Automated Deduction (CADE-11)*, Saratoga Springs, New York.

- Fegaras, L. & Stemple, D., 1991. Using Type Transformation in Database System Implementation. In *3rd International Conference on Database Programming Languages*, Kanelakis, P. & Schmidt, J.W. (ed.), Morgan Kaufmann, Proc. 3rd International Workshop on Database Programming Languages, Nafplion, Greece pp 337-353.
- Fenichel, R.R. & Yochelson, J.C., 1969. A LISP Garbage Collector for Virtual-Memory Systems. *Communications of the ACM* 12, 11 pp 611-612.
- Ghelli, G., Orsini, R., Paz, A.P. & Trinder, P., 1992. Design of an Integrated Query and Manipulation Notation for Database Languages . ESPRIT BRA Project 3070 FIDE Technical Report FIDE/92/41.
- Goldberg, A. & Robson, D., 1983. *Smalltalk-80: The Language and its Implementation*. Addison Wesley, Reading, Massachusetts.
- Greenwood, R.M., Guy, M.R. & Robinson, D.J.K., 1992. The Use of a Persistent Language in the Implementation of a Process Support System. *ICL Technical Journal* 8, 1 pp 108-130.
- Griswold, R.E., Poage, J.F. & Polonsky, I.P., 1971. *The SNOBOL4 Programming Language*. Prentice-Hall, Englewood Cliffs, New Jersey.
- Grossman, R.L., Qin, X., Xu, W., Ramamoorthy, H. & Araujo, N., 1994. Managing Physical Folios of Objects Between Nodes. In *Persistent Object Systems, Tarascon 1994*, Atkinson, M.P., Maier, D. & Benzaken, V. (ed.), Springer-Verlag, Proc. 6th International Workshop on Persistent Object Systems, Tarascon, France pp 217-234.
- Gruber, O., 1992. Eos, an Environment for Persistent and Distributed Applications over a Shared Object Space. PhD Thesis, University of Université Pierre et Marie Curie, Paris VI.
- Gruber, O., Amsaleg, L., Daynès, L. & Valduriez, P., 1992. Eos, an Environment for Object-based Systems. In *Proc. 25th International Conference on Systems Sciences*, Hawaii pp 757-768.
- Hammer, M. & McLeod, D., 1981. Database Description with SDM: A Semantic Database Model. *ACM Transactions on Database Systems* 6, 3 pp 351-386.
- Han, J. & Welsh, J., 1993. Methodology Modelling: Combining Software Processes with Software Products . University of Queensland Software Validation Research Centre Technical Report 93-17.
- Hoare, C.A.R., 1974. Monitors: An Operating System Structuring Concept. *Communications of the ACM* 17, 10 pp 549-557.
- Hughes, J.G. & Connolly, M., 1990. Data Abstraction and Transaction Processing in the Database Programming Language RAPP. In *Advances in Database Programming Languages*, Bancilhon, F. & Buneman, O.P. (ed.), Addison-Wesley and ACM Press pp 177-186.
- Hurst, A.J. & Sajejev, A.S.M., 1990. A Capability Based Language for Persistent Programming. In *Persistent Object Systems*, Rosenberg, J. & Koch, D.M. (ed.), Springer-Verlag pp 186-201.
- Kanellakis, P. & Schmidt, J.W. (ed.), 1991. *Database Programming Languages: Bulk Types and Persistent Data*. Morgan Kaufmann.

- Kemper, A. & Kossmann, D., 1995. Adaptable Pointer Swizzling Strategies in Object Bases: Design, Realization and Quantitative Analysis. *VLDB Journal* 4, 3.
- King, F., 1978. IBM Report on the Contents of a Sample of Programs Surveyed . IBM, San Jose, California.
- Kirby, G.N.C., 1992a. Persistent Programming with Strongly Typed Linguistic Reflection. In *Proc. 25th International Conference on Systems Sciences*, Hawaii pp 820-831.
- Kirby, G.N.C., 1992b. Reflection and Hyper-Programming in Persistent Programming Systems. Ph.D. Thesis, University of St Andrews.
- Kirby, G.N.C., Brown, A.L., Connor, R.C.H., Cutts, Q.I., Dearle, A., Moore, V.S., Morrison, R. & Munro, D.S., 1994. The Napier88 Standard Library Reference Manual Version 2.2 . University of St Andrews Technical Report CS/94/7.
- Kirby, G.N.C., Connor, R.C.H., Cutts, Q.I., Dearle, A., Farkas, A.M. & Morrison, R., 1992. Persistent Hyper-Programs. In *Persistent Object Systems*, Albano, A. & Morrison, R. (ed.), Springer-Verlag, Proc. 5th International Workshop on Persistent Object Systems, San Miniato, Italy pp 86-106.
- Kirby, G.N.C., Connor, R.C.H. & Morrison, R., 1994. START: A Linguistic Reflection Tool Using Hyper-Program Technology. In *Persistent Object Systems, Tarascon 1994*, Atkinson, M.P., Maier, D. & Benzaken, V. (ed.), Springer-Verlag, Proc. 6th International Workshop on Persistent Object Systems, Tarascon, France pp 355-373.
- Kirby, G.N.C. & Dearle, A., 1990. An Adaptive Graphical Browser for Napier88 . University of St Andrews Technical Report CS/90/16.
- Koch, B., Schunke, T., Dearle, A., Vaughan, F., Marlin, C., Fazakerley & Barter, C., 1990. Cache Coherence and Storage Management in a Persistent Object System. In *Implementing Persistent Object Bases*, Dearle, A., Shaw, G. & Zdonik, S.B. (ed.), Morgan Kaufmann pp 103-113.
- Koch, J., Mall, M., Putfarken, P., Reimer, M., Schmidt, J.W. & Zehnder, C.A., 1983. Modula/R Report Lilith Version . ETH Zürich.
- Kolodner, E., 1987. Recovery Using Virtual Memory. M.Sc. Thesis, University of MIT.
- Kolodner, E., Liskov, B. & Weihl, W., 1989. Atomic Garbage Collection: Managing a Stable Heap. In *Proc. ACM SIGMOD International Conference on the Management of Data*, Portland, Oregon pp 15-25.
- Krablin, G.L., 1987a. Building Flexible Multilevel Transactions in a Distributed Persistent Environment. In *Proc. 2nd International Workshop on Persistent Object Systems*, Appin, Scotland pp 86-117.
- Krablin, G.L., 1987b. Transactions and Concurrency . Universities of Glasgow and St Andrews Technical Report PPRR-46-87.
- Kulkarni, K.G., 1994. ADT-based Type System for SQL. In *Query Processing for Advanced Database Systems*, Freytag, J.C., Maier, D. & Vossen, G. (ed.), Morgan Kaufmann.

- Kung, H.T. & Robinson, J.T., 1982. On Optimistic Methods for Concurrency Control. *ACM Transactions on Database Systems* 6, 2 pp 213-226.
- Lavery, D.O., 1995a. The Design of Effective Software Visualizations for Persistent Programming Languages . ESPRIT BRA Project 6309 FIDE₂ Technical Report FIDE/95/116.
- Lavery, D.O., 1995b. Towards Visualizing Persistent Stores. *in preparation*.
- Levin, R. & McJones, P.R., 1993. The Vesta Approach to Precise Configuration of Large Software Systems . DEC Systems Research Center Technical Report 105.
- Liskov, B., 1988. Distributed Programming in Argus. *Communications of the ACM* 31, 3 pp 300-312.
- Liskov, B., Johnson, P., Gruber, R. & Shriram, L., 1990. A Highly Available Object Repository for Use in a Heterogeneous Distributed System. In *Implementing Persistent Object Bases, Principles and Practice*, Dearle, A., Shaw, G.M. & Zdonik, S.B. (ed.), Morgan Kaufmann, Proc. 4th International Workshop on Persistent Object Systems, Martha's Vineyard, USA pp 255-266.
- Matthes, F., 1992. Generic Database Programming: A Linguistic and Architectural Framework. PhD Thesis, University of Hamburg.
- Matthes, F., Müller, R. & Schmidt, J.W., 1992. Object Stores as Servers in Persistent Programming Environments—The P-Quest Experience . ESPRIT BRA Project 3070 FIDE Technical Report FIDE/92/48.
- Matthes, F. & Müßig, S., 1993. The Tycoon Language TL: An Introduction . University of Hamburg Technical Report DBIS 112-93.
- Matthes, F., Müßig, S. & Schmidt, J.W., 1994. Persistent Polymorphic Programming in Tycoon: An Introduction . ESPRIT BRA Project 6309 FIDE₂ Technical Report FIDE/94/106.
- Matthes, F. & Schmidt, J.W., 1989. The Type System of DBPL. In *Proc. 2nd International Workshop on Database Programming Languages*, Salishan, Oregon pp 219-225.
- Matthes, F. & Schmidt, J.W., 1991. Bulk Types: Built-In or Add-On? In *3rd International Workshop on Database Programming Languages*, Kanelakis, P. & Schmidt, J.W. (ed.), Morgan Kaufmann, Proc. 3rd International Workshop on Database Programming Languages, Nafplion, Greece pp 33-54.
- Matthes, F. & Schmidt, J.W., 1992. Definition of the Tycoon Language TL - A Preliminary Report . University of Hamburg, Germany Technical Report FBI-HH-B-160/92.
- Matthes, F. & Schmidt, J.W., 1994. Persistent Threads. In *Proc. 20th International Conference on Very Large Data Bases*, Santiago, Chile pp 403-414.
- Matthews, D.C.J., 1982. Poly Report . University of Cambridge Technical Report 28.
- Matthews, D.C.J., 1985. A Persistent Storage System for Poly . Cambridge University.
- Matthews, D.C.J., 1989. Papers on Poly/ML . University of Cambridge.

- McCarthy, J., Abrahams, P.W., Edwards, D.J., Hart, T.P. & Levin, M.I., 1962. *The Lisp Programmers' Manual*. M.I.T. Press, Cambridge, Massachusetts.
- Microsoft Corporation, 1994a. Microsoft Access: Building Applications . .
- Microsoft Corporation, 1994b. Microsoft Access: User's Guide . .
- Milner, R., 1978. A Theory of Type Polymorphism in Programming. *Journal of Computer and System Sciences* 17, 3 pp 348-375.
- Milner, R., 1980. *Lecture Notes in Computer Science Vol 92: A Calculus of Communicating Systems*. Springer-Verlag.
- Milner, R., 1991. The Polyadic π -Calculus: A Tutorial . University of Edinburgh Technical Report ECS-LFCS-92-249.
- Mira da Silva, M., Atkinson, M.P. & Black, A., 1995. Semantics for Parameter Passing in a Type-complete Persistent RPC. *Submitted to: 6th International Workshop on Database Programming Languages, Gubbio, Italy*.
- Mira da Silva, M.M., 1995. Automating Type-safe RPC. *In Proc. 5th International Workshop on Research Issues on Data Engineering: Distributed Object Management, Taipei, Taiwan* pp 100-107.
- Morrison, R., 1979. On the Development of Algol. Ph.D. Thesis, University of St Andrews.
- Morrison, R. & Atkinson, M.P., 1992. Persistent Object Systems. *In Proc. 25th International Conference on Systems Sciences, Hawaii* pp 754-755.
- Morrison, R., Atkinson, M.P. & Dearle, A., 1987. Flexible Incremental Bindings in a Persistent Object Store . Universities of Glasgow and St Andrews Technical Report PPRR-38-87.
- Morrison, R., Baker, C., Connor, R.C.H., Cutts, Q.I. & Kirby, G.N.C., 1993. Approaching Integration in Software Environments . University of St Andrews Technical Report CS/93/10.
- Morrison, R., Barter, C.J., Brown, A.L., Carrick, R., Connor, R.C.H., Dearle, A., Hurst, A.J. & Livesey, M.J., 1989. Language Design Issues in Supporting Process-Oriented Computation in Persistent Environments. *In Proc. 22nd International Conference on System Sciences, Hawaii* pp 736-744.
- Morrison, R., Barter, C.J., Connor, R.C.H., Denton, J., Kirkpatrick, G., Munro, D.S., Pretsell, B. & Stemple, D., 1993. Concurrency Control in Process Models. *IOPENER* 2, 1 pp 11-12.
- Morrison, R., Brown, A.L., Carrick, R., Connor, R.C.H. & Dearle, A., 1988. On the Integration of Object-Oriented and Process-Oriented Computation in Persistent Environments. *In Lecture Notes in Computer Science 334*, Dittrich, K.R. (ed.), Springer-Verlag pp 334-339.
- Morrison, R., Brown, A.L., Carrick, R., Connor, R.C.H., Dearle, A. & Atkinson, M.P., 1987. Polymorphism, Persistence and Software Reuse in a Strongly Typed Object Oriented Environment. *Software Engineering Journal*, December pp 199-204.

- Morrison, R., Brown, A.L., Carrick, R., Connor, R.C.H., Dearle, A. & Atkinson, M.P., 1990. The Napier Type System. In *Persistent Object Systems*, Rosenberg, J. & Koch, D.M. (ed.), Springer-Verlag, Proc. 3rd International Workshop on Persistent Object Systems, Newcastle, Australia pp 3-18.
- Morrison, R., Brown, A.L., Connor, R.C.H., Cutts, Q.I., Dearle, A., Kirby, G.N.C. & Munro, D.S., 1994. The Napier88 Reference Manual (Release 2.0) . University of St Andrews Technical Report CS/94/8.
- Morrison, R., Brown, A.L., Connor, R.C.H., Cutts, Q.I., Kirby, G.N.C., Dearle, A., Rosenberg, J. & Stemple, D., 1990. Protection in Persistent Object Systems. In *Security and Persistence*, Rosenberg, J. & Keedy, J.L. (ed.), Springer-Verlag, Proc. International Workshop on Security and Persistence, Bremen, 1990 pp 48-66.
- Morrison, R., Brown, A.L., Connor, R.C.H. & Dearle, A., 1989. The Napier88 Reference Manual . Universities of Glasgow and St Andrews Technical Report PPRR-77-89.
- Morrison, R., Brown, A.L., Dearle, A. & Atkinson, M.P., 1986. An Integrated Graphics Programming Environment. *Computer Graphics Forum* 5, 2 pp 147-157.
- Morrison, R., Brown, A.L., Dearle, A. & Atkinson, M.P., 1990. On the Classification of Binding Mechanisms. *Information Processing Letters* 34 pp 51-55.
- Morrison, R., Connor, R.C.H., Cutts, Q.I., Dunstan, V.S. & Kirby, G.N.C., 1995. Exploiting Persistent Linkage in Software Engineering Environments. *Accepted By: Computer Journal*.
- Morrison, R., Connor, R.C.H., Cutts, Q.I. & Kirby, G.N.C., 1994. Persistent Possibilities for Software Environments. In *The Intersection between Databases and Software Engineering*, IEEE Computer Society Press, Proc. ICSE-16 Workshop on the Intersection between Databases and Software Engineering, Sorrento, Italy pp 78-87.
- Morrison, R., Dearle, A., Bailey, P.J., Brown, A.L. & Atkinson, M.P., 1985. The Persistent Store as an Enabling Technology for Integrated Project Support Environments. In *Proc. 8th IEEE International Conference on Software Engineering*, London pp 166-172.
- Morrison, R., Dearle, A., Connor, R.C.H. & Brown, A.L., 1991. An Ad-Hoc Approach to the Implementation of Polymorphism. *ACM Transactions on Programming Languages and Systems* 13, 3 pp 342-371.
- Moss, E. & Sinofsky, A., 1988. Managing Persistent Data with Mnome: Designing a Reliable, Shared Object Interface. In *Advances in Object-Oriented Database Systems*, Springer-Verlag pp 298-316.
- Moss, J.E.B., 1990. Working with Persistent Objects: To Swizzle or Not to Swizzle . COINS, University of Massachusetts Technical Report 90-38.
- Moss, J.E.B. & Hosking, A.L., 1994. Expressing Object Residency Optimizations using Pointer Type Annotations. In *Persistent Object Systems, Tarascon 1994*, Atkinson, M.P., Maier, D. & Benzaken, V. (ed.), Springer-Verlag, Proc. 6th International Workshop on Persistent Object Systems, Tarascon, France pp 3-15.

- Munro, D.S., 1993. On the Integration of Concurrency, Distribution and Persistence. Ph.D. Thesis, University of St Andrews.
- Munro, D.S., Connor, R.C.H., Morrison, R., Scheuerl, S. & Stemple, D., 1994. Concurrent Shadow Paging in the Flask Architecture. In *Persistent Object Systems, Tarascon 1994*, Atkinson, M.P., Maier, D. & Benzaken, V. (ed.), Springer-Verlag, Proc. 6th International Workshop on Persistent Object Systems, Tarascon, France pp 16-42.
- Mylopoulos, J., Bernstein, P.A. & Wong, H.K.T., 1980. A Language Facility for Designing Database-Intensive Applications. *Association for Computing Machinery Transactions on Database Systems* 5, 2 pp 185-207.
- Nettles, S. & Wing, J., 1992. Persistence + Undoability = Transactions. In *Proc. 25th International Conference on Systems Sciences*, Hawaii pp 832-843.
- Nodine, M.H. & Zdonik, S.B., 1992. Co-operative Transaction Hierarchies: Transaction Support for Design Applications. *VLDB Journal* 1, 1 pp 41-80.
- Ohuri, A., Buneman, P. & Breazu-Tannen, V., 1989. Database Programming in Machiavelli - a Polymorphic Language with Static Type Inference. In *Proc. SIGMOD 89* pp 46-57.
- Organick, E.I., 1972. *The Multics System: An Examination of its Structure*. M.I.T. Press, Cambridge, Mass.
- Powell, M.S., 1985. Adding Programming Facilities to an Abstract Data Store. In *Proc. 1st International Workshop on Persistent Object Systems*, Appin, Scotland pp 139-160.
- PS-algol, 1988. PS-algol Reference Manual, 4th edition . Universities of Glasgow and St Andrews Technical Report PPRR-12-88.
- Rees, J. & Clinger, W., 1986. Revised Report on the Algorithmic Language Scheme. *ACM SIGPLAN Notices* 21, 12 pp 37-43.
- Reinwald, B., Dessloch, S., Carey, M., Lehman, T., Pirahesh, H. & Srinivasan, V., 1994. Making Real Data Persistent: Initial Experiences with SMRC. In *Persistent Object Systems, Tarascon 1994*, Atkinson, M.P., Maier, D. & Benzaken, V. (ed.), Springer-Verlag, Proc. 6th International Workshop on Persistent Object Systems, Tarascon, France pp 202-216.
- Richards, M. & Whitby-Strevans, C., 1979. *BCPL—The Language and its Compiler*. Cambridge University Press, Cambridge, UK.
- Richardson, J. & Carey, M., 1989. Persistence in the E Language: Issues and Implementation. *SPE* 19, 12 pp 1115 - 1150.
- Richardson, J.E. & Carey, M.J., 1990. Implementing Persistence in E. In *Persistent Object Systems*, Rosenberg, J. & Koch, D.M. (ed.), Springer-Verlag pp 175-199.
- Rosenberg, J., 1990. The MONADS Architecture - A Layered View. In *Implementing Persistent Object Bases*, Dearle, A., Shaw, G.M. & Zdonik, S.B. (ed.), Morgan Kaufmann, Proc. 4th International Workshop on Persistent Object Systems, Martha's Vineyard, USA pp 215-225.

- Rosenberg, J. & Dearle, A., 1995. Proceedings of Minitrack on Distribution and Concurrency in Persistent Systems. *In Proc. 28th International Conference on Systems Sciences*, Hawaii.
- Rosenberg, J., Henskens, F., Brown, A.L., Morrison, R. & Munro, D., 1990. Stability in a Persistent Store Based on a Large Virtual Memory. *In Security and Persistence*, Rosenberg, J. & Keedy, J.L. (ed.), Springer-Verlag, Proc. International Workshop on Security and Persistence, Bremen, 1990 pp 229-245.
- Rosenberg, J. & Keedy, J.L. (ed.), 1990. *Security and Persistence*. Springer-Verlag.
- Ruffin, M., 1992. Kitlog, A Generic Logging Service. *In Proc. 11th IEEE Symposium on Reliable Distributed Systems*, Houston, TX pp 139-146.
- Russell, G., 1994. DOLPHIN: Persistent, Object-oriented and Networked. PhD Thesis, University of University of Strathclyde.
- Russell, G., Shaw, P. & Cockshott, W.P., 1994. DAIS: An Object-Addressed Processor Cache. *In Persistent Object Systems, Tarascon 1994*, Atkinson, M.P., Maier, D. & Benzaken, V. (ed.), Springer-Verlag, Proc. 6th International Workshop on Persistent Object Systems, Tarascon, France pp 374-386.
- Schaffert, C., Cooper, T. & Wilpot, C., 1985. Trellis Object-Based Environment Language Reference Manual . DEC Systems Research Center.
- Scheuerl, S.J.G., Connor, R.C.H., Morrison, R., Moss, J.E.B. & Munro, D.S., 1995. MaStA - An I/O Cost Model for Database Crash Recovery Mechanisms . University of St Andrews Technical Report CS/95/1.
- Schmidt, J.W., 1977. Some High Level Language Constructs for Data of Type Relation. *ACM Transactions on Database Systems* 2, 3 pp 247-261.
- Schmidt, J.W. & Matthes, F., 1992. The Database Programming Language DBPL Rationale and Report . ESPRIT BRA Project 3070 FIDE Technical Report FIDE/92/46.
- ServioLogic Ltd, 1987. Programming in OPAL . .
- Shapiro, M., Gautron, P. & Mosseri, L., 1989. Persistence and Migration for C++ Objects. *In Proc. Proceedings, European Conference on Object-Oriented Programming (ECOOP)*.
- Sheard, T., 1991. Automatic Generation and Use of Abstract Structure Operators. *ACM Transactions on Programming Languages and Systems* 19, 4 pp 531-557.
- Sheard, T. & Hook, J., 1994. Type Safe Meta-Programming . Oregon Graduate Institute.
- Sheard, T. & Stemple, D., 1989. Automatic Verification of Database Transaction Safety. *ACM Transactions on Database Systems* 12, 3 pp 322-368.
- Shipman, D., 1981. The Functional Data Model and the Data Language DAPLEX. *Association for Computing Machinery Transactions on Database Systems* 6, 1 pp 140-173.

- Singhal, V., Kakkad, S.V. & Wilson, P.R., 1992. Texas: An Efficient, Portable Persistent Store. In *Persistent Object Systems*, Albano, A. & Morrison, R. (ed.), Springer-Verlag, Proc. 5th International Workshop on Persistent Object Systems, San Miniato, Italy pp 11-33.
- Sjøberg, D., 1991. The Thesaurus – A Tool for Meta Data Management . ESPRIT BRA Project 3070 FIDE Technical Report FIDE/91/6.
- Sjøberg, D.I.K., 1993. Thesaurus-Based Methodologies and Tools for Maintaining Persistent Application Systems. Ph.D. Thesis, University of Glasgow.
- Sjøberg, D.I.K., Atkinson, M.P., Lopes, J.C. & Trinder, P.W., 1993. Building an Integrated Persistent Application. In *Database Programming Languages*, Beerl, C., Ohori, A. & Shasha, D.E. (ed.), Springer-Verlag, Proc. 4th International Conference on Database Programming Languages, New York City pp 359-375.
- Sjøberg, D.I.K., Atkinson, M.P. & Welland, R.C., 1994a. Thesaurus-Based Software Environments. In *The Intersection between Databases and Software Engineering*, IEEE Computer Society Press, Proc. ICSE-16 Workshop on the Intersection between Databases and Software Engineering, Sorrento, Italy pp 41-46.
- Sjøberg, D.I.K., Cutts, Q.I., Welland, R. & Atkinson, M.P., 1994b. Analysing Persistent Language Applications. In *Persistent Object Systems, Tarascon 1994*, Atkinson, M.P., Maier, D. & Benzaken, V. (ed.), Springer-Verlag, Proc. 6th International Workshop on Persistent Object Systems, Tarascon, France pp 235-255.
- Sjøberg, D.I.K., Philbrow, P.C., Waite, C. & Welland, R., 1995. Build Management in Database Programming Language Environments. *Submitted to: 6th International Workshop on Database Programming Languages, Gubbio, Italy.*
- Smith, J.M., Fox, S. & Landers, T., 1981. Reference Manual for ADAPLEX . Computer Corporation of America Technical Report CCA-81-02.
- Sommerville, I., Welland, R., Potter, S. & Smart, J., 1989. The ECLIPSE User Interface. *Software – Practice and Experience* 19, 4 pp 371.
- Stanley, M., 1986. An Evaluation of the Flex PSE . Defence Research Agency, Malvern, England Technical Report 86003.
- Stanley, M. & Drummond, P.D., 1988. A Flexible Basis for Software Configuration Management . Defence Research Agency, Malvern, England Technical Report 4127.
- Stemple, D., Fegaras, L., Sheard, T. & Socorro, A., 1990. Exceeding the Limits of Polymorphism in Database Programming Languages. In *Lecture Notes in Computer Science 416*, Bancilhon, F., Thanos, C. & Tsichritzis, D. (ed.), Springer-Verlag pp 269-285.
- Stemple, D. & Morrison, R., 1992. Specifying Flexible Concurrency Control Schemes: An Abstract Operational Approach. In *Proc. 15th Australian Computer Science Conference*, Hobart, Tasmania pp 873-891.
- Stemple, D., Sheard, T. & Fegaras, L., 1992a. Linguistic Reflection: A Bridge from Programming to Database Languages. In *Proc. 25th International Conference on Systems Sciences*, Hawaii pp 844-855.

- Stemple, D., Stanton, R.B., Sheard, T., Philbrow, P., Morrison, R., Kirby, G.N.C., Fegaras, L., Cooper, R.L., Connor, R.C.H., Atkinson, M.P. & Alagic, S., 1992b. Type-Safe Linguistic Reflection: A Generator Technology . ESPRIT BRA Project 3070 FIDE Technical Report FIDE/92/49.
- Stonebraker, M. & Kemnitz, G., 1991. The POSTGRES Next Generation Database Management System. *Communications of the ACM* 34, 10 pp 78-92.
- Strachey, C., 1967. *Fundamental Concepts in Programming Languages*. Oxford University Press, Oxford.
- Straw, A.F., Mellender, F. & Riegel, S., 1984. Object Management in a Persistent Smalltalk System. *SPE* 19, 8 pp 719-737.
- Sutton, S., 1990. A Flexible Consistency Model for Persistent Data in Software-Process Programming. In *Implementing Persistent Object Bases, Principles and Practice*, Dearle, A., Shaw, G.M. & Zdonik, S.B. (ed.), Morgan Kaufmann, Proc. 4th International Workshop on Persistent Object Systems, Martha's Vineyard, USA pp 305-319.
- Suzuki, S., Kitsuregawa, M. & Takagi, M., 1994. An Efficient Pointer Swizzling Method for Navigation Intensive Applications. In *Persistent Object Systems, Tarascon 1994*, Atkinson, M.P., Maier, D. & Benzaken, V. (ed.), Springer-Verlag, Proc. 6th International Workshop on Persistent Object Systems, Tarascon, France pp 79-95.
- Teitelbaum, T. & Reps, T., 1981. The Cornell Program Synthesizer: A Syntax-Directed Programming Environment. *Communications of the ACM* 24, 9 pp 563-573.
- Tennent, R.D., 1977. Language Design Methods Based on Semantic Principles. *Acta Informatica* 8 pp 97-112.
- Thatte, S.M., 1986. Persistent Memory: A Storage Architecture for Object Oriented Database Systems. In *Proc. ACM/IEEE International Workshop on Object-Oriented Database Systems*, Pacific Grove, California pp 148-159.
- Thomas, I., 1989. PCTE Interfaces: Supporting Tools in Software-Engineering Environments. *IEEE Software* pp 15-23.
- Trinder, P.W. & Wadler, P.L., 1989. Improving List Comprehension Database Queries. In *Proc. TENCON'89*, Bombay, India pp 186-192.
- van Wijngaarden, A., Mailloux, B.J., Peck, J.E.L. & Koster, C.H.A., 1969. Report on the Algorithmic Language ALGOL 68. *Numerische Mathematik* 14 pp 79-218.
- Wai, F., 1989. Distributed Concurrent Persistent Languages: an Experimental Design and Implementation . Universities of Glasgow and St Andrews Technical Report PPRR-76-89.
- Waite, C., 1995. The Glasgow Persistent Workshop: User Documentation. *in preparation*.
- Wasserman, A.I., Shertz, D.D., Kersten, M.L., Reit, R.P. & van de Dippe, M.D., 1981. Revised Report on the Programming Language PLAIN. *ACM SIGPLAN Notices* 5, 16.

- Wetzel, I., 1994. Programming with STYLE: On the Systematic Development of Programming Environments. PhD Thesis, University of Hamburg.
- Wileden, J., Wolf, A., Fisher, C. & Tarr, P., 1988. PGRAPHITE: An Experiment in Persistent Typed Object Management. *In Proc. ACM SIGSOFT '88: Third Symposium on Software Development Environments* pp 130-142.
- Wilson, P.R., 1990. Pointer Swizzling at Page Fault Time: Efficiently Supporting Huge Address Spaces on Standard Hardware . University of Illinois at Chicago Technical Report UIC-EECS-90-6.
- Wilson, P.R., 1992. Uniprocessor Garbage Collection Techniques. In *Lecture Notes in Computer Science 637*, Springer-Verlag, Proc. International Workshop on Memory Management (St Malo, France, September 1992) pp 1-42.
- Wirth, N., 1971. The Programming Language Pascal. *Acta Informatica* 1 pp 35-63.
- Zeula, P. & Rabitti, F., 1992. Navigation Index for an Object Store. *In Proc. 25th International Conference on System Sciences*, Hawaii pp 756-765.