

University of Glasgow
Department of Computing Science

Lilybank Gardens
Glasgow G12 8QQ



University of St. Andrews
Department of Computational Science



North Haugh
St Andrews KY16 8SX

A Persistent Information
Space Architecture

Persistent Programming
Research Report 21

Alan Dearle.

Preface

This report is a version of a paper which has been submitted to the
IFIP'86 Conference, Dublin, August, 1986.

A persistent information space architecture

Malcolm P. Atkinson, Ronald Morrison and Graham D. Pratten

University of Glasgow, Glasgow, Scotland G12 8QQ. Tel. 0413398855,

University of St Andrews, St Andrews, Scotland KY16 9SX. Tel. 033476161

and

International Computers Ltd., Kidsgrove, England. ST7 1TL. Tel. 078229681.

Abstract

Present day computer systems depend on a plethora of mechanisms that by their lack of coherence increase the cost of developing applications. By integrating these mechanisms we reduce the complexity of the system and thereby achieve savings throughout product life cycles. We identify five major areas of system design for modern systems: controlling complexity, orthogonal persistence, controlled system evolution, protection of data and concurrent computation. We report on our current research on programming languages and environments and propose a persistent information space architecture (PISA) capable of integrating all these activities.

1. Introduction

If we examine existing computer systems we find many dichotomies and discontinuities in their design and in the design of their user interfaces. Some of these were deliberately introduced during the 1960's and 1970's to satisfy reasonable performance constraints. Others simply reflect the order in which the basic design concepts of computer systems have emerged over a period of twenty to thirty years. We contend that the present dependence on this plethora of mechanisms such as command languages, editors, file systems, compilers and interpreters, linkage editors and binders, debuggers, DBMS sublanguages, graphics libraries increases the cost of understanding and maintaining software, and training programmers, for even the simplest of activities. It is important to remove this incoherence now since it is placing a considerable overhead on the users and developers of computer systems, hindering their full exploitation, and delaying the development of integrated projects support environments (IPSEs) and intelligent knowledge based systems (IKBSs). The removal of this incoherence is now possible because of the dramatic shift in the cost of hardware relative to software. In this paper we report on our current research on programming languages and environments and propose a persistent information space architecture (PISA) capable of integrating all programming activities.

Inherent in our approach is that any program must be integrated with its environment. We have identified the following problems faced by users of current systems which when negated become requirements of modern systems in order to achieve simplicity and integration. We do not regard this list as exhaustive. Indeed we would expect that having achieved these requirements new ones will emerge in the pursuit of better programming systems. They are

- a. **Controlling complexity:** The complexity of the system must be kept under control, so that developers and users can concentrate on the application rather than the complexity of the system. This depends on establishing consistent rules which apply throughout the design and being parsimonious in the introduction of new concepts into these designs.
- b. **Orthogonal persistence:** The discontinuity between the method of using data that is short term and manipulated by program and long term data that is manipulated by

the file system or DBMS causes unnecessary complexity. We have defined the persistence of data to be the length of time for which the data exists and is useable[2]. We aspire to systems where the use of data is independent of its persistence.

- c. **Controlled system evolution:** The uses of data (including program) are neither limited nor predictable. It is necessary to support the construction of unanticipated software systems or databases which make use of pre-existing data (or program) even when the data and program were defined independently of one another. For large scale, widely used or continuously used systems any alteration to part of the system should not require total rebuilding. We require a mechanism which will allow the programmer to control the units of reconstruction.
- d. **Protection of data:** Large bodies of data are inherently valuable. It is necessary to protect them from misuse and from hardware and software failure. This implies both a type and protection system to meet all our needs and recovery mechanisms to limit the losses due to component failure.
- e. **Concurrent computation:** A large body of data requires a community effort for its construction and maintenance. Any useful body of data is likely to be of concurrent interest to many users, probably in dispersed geographic locations. Different models of concurrency and transactions may have to be accommodated by the underlying mechanism.

We have designed and implemented the language PS-algo [3] as a testbed for experiments on the above requirements. We report here on some results and propose further experiments in the search for better programming systems.

2. Controlling complexity

In many ways the main task of computer science is to invent and implement appropriate languages to suit applications at any level, both hardware and software, in order that programming may become easier and cheaper. We may therefore use the experience of programming language designers in controlling complexity with the aim of building a total system that will integrate all the naming, binding and other mechanisms of programming languages, database systems and operating

systems.

Central to our method of language design is that the language itself should be simple. Merely adding features to languages without integrating them into some overall design simply increases the complexity of the system and can often lead to it being beyond the intellectual capacity of the programmer and sometimes even the implementor. This complexity is often due, in part at least, to being too restrictive but can be avoided by the language designer formulating the fundamental concepts behind the language and generalising these wherever possible. Simplicity is achieved by having no exceptions to these general rules.

Our languages are designed using three principles due to Landin[16] and Strachey[21] and further developed by Tennent[22]. They are

1. The principle of correspondence: the use of names should be consistent within a language. In particular there should be a one to one correspondence between the method of introducing names in declarations and in parameter lists.
2. The principle of abstraction: all the major syntactic categories should have abstractions defined over them. For example, functions are abstractions over expressions.
3. The principle of data type completeness: all data types should be first class without arbitrary restriction on their use.

We have described our design methodology fully elsewhere[12] and will not labour it here. However in applying these rules no exceptions are allowed since having exceptions makes the language more complex in terms of defining rules and usually less powerful since there are restrictions. Designing languages using the above rules yields languages that are free from inconsistencies and idiosyncrasies in this manner. Thus the rules can be used as a metric for comparing languages. It is part of our research into programming language design to find further rules, if there are any, to aid the design process. We have applied this general philosophy to the design of programming languages and programming environments in the PS-algol system.

3. Orthogonal Persistence

We seek to eliminate the differences between programming language models of data and the environment models of data. This can be done by separating the issue of what data structures are best

for a particular problem from the issue of identifying and preserving the data. This second property we call persistence. A spectrum of persistence exists and is categorised by

1. transient results in expression evaluation.
2. local variables in procedure activations.
3. own variables, global variables and heap items whose extent is different from their scope.
4. data that exists between executions of a program.
5. data that exists between various versions of a program.
6. data that outlives the program.

The first three persistence categories are usually supported by programming languages and the second three categories by a DBMS, whereas filing systems are predominantly used for categories 4 and

5. Persistence is therefore an abstraction over the programming language and file system or DBMS views of data.

There are three major advantages in having only one model of data. Firstly in any traditional program there is usually a considerable amount of code, typically 30% of the total, concerned with transferring data to and from files or a DBMS. Much space and time is taken up by code to perform translations between the program's form of data and the form used for the long term storage medium. For example, we normally have to explicitly flatten and rebuild graphs or trees modelled in the programming language in order to write them out or read them back in from the file store. If this activity is performed by the underlying system then every application program would be 30% smaller in terms of code, with all the attendant benefits that entails throughout the life cycle of the product.

The second major advantage is that it is conceptually simpler to have only one program view of data. Again in traditional systems there are three models of data, one in the real world, one in the database system and one in the program's runtime system. This is unsatisfactory since the programmer or system designer has to visualise and maintain all of these models correctly. This leads to an intellectual overhead in using these models and preventing them from becoming mutually inconsistent.

The third major advantage is that the data type protection offered by the programming language on its data is now available for all data. Thus the main protection mechanism is not lost across an unnecessary mapping.

We regard the provision of persistence as orthogonal to all other properties of data. For this we define the Principle of Persistence Independence as

The persistence of a data object is independent of how the program manipulates that data object.

That is, all code should be written so that it will work with the same interpretation independent of the persistence of the data on which it operates. This reduces the number of conceptual mappings from three to one - the one between the real world and the program model and thus simplifies the task of building an application system.

We also identify an extension to the Principle of Data Type Completeness to handle persistence. It is

All data objects should be allowed the full range of persistence.

Large systems need long lived data since they all need to store their code and data. Computing systems that involve humans must support long term activity. Since the attention span of the human is short in computing terms then the storage of the present state of an attempted solution to a problem is essential. To achieve this in an integrated manner the system should provide persistence designed by the above principles.

4. Controlled System Evolution

In traditional programming languages, operating systems and file systems there are a number of binding mechanisms which are often not easy to comprehend or use. A binding mechanism has four components by which it can be categorised. They are

- a. what does the name bind to?
- b. when is the binding performed?
- c. what scoping is involved?
- d. when is the type checking performed?

Names of variables bind to locations whose value may change without altering the binding.

Constant names bind to values. The binding of both variables and constants is usually performed when the location is created or the value calculated. That is, at run time unless of course the location or value is manifest (a literal) when the binding may be performed at compile time. Manifest constants can be seen in BCPL[20] or Pascal[24] and manifest locations are the variables of Fortran. In block structured languages it is usual for variables to bind to locations created at run time. Constants whose values are created dynamically can be seen in S-algo[12], the simple values of Ada[15] and most applicative languages.

Names may be scoped statically in their compile time environment or dynamically in their run time environment. The traditional algo[1] scoping rule is static whereas dynamic scoping can be seen in Lisp[18] or the segment binding mechanism of Multics[11].

Type checking may be performed statically by the compiler or dynamically 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 in read statements and in projections out of a union. Some languages such as SASL[23] deliberately choose run time type checking to facilitate polymorphism. For example it is difficult to write down or deduce the type of the self apply function

$\lambda x. x x$

but simple to check it at run time.

There are potentially 16 different methods of binding commonly in use in modern computer systems based on the four binding choices given above. The most static form of binding is where only manifest constants are allowed with static scoping and static type checking. The most dynamic form allows variables with non manifest values, dynamic scoping and dynamic type checking. We coin the term flexible incremental binding (FIB) to describe this mixture of bindings, which we expect to obey the Principle of Correspondence. In practice most languages have more than one binding mechanism.

The example of Pascal is interesting here where we observe that variable names may be bound to locations but not manifest locations and constant names may be bound to manifest values but not values created dynamically. The scoping rule is static except for file names where it is evaluated in the environment and is therefore dynamic. Type checking is static except in read statements and also on the projection out of a union of variants both of which require a dynamic check.

4.1 Name Spaces

The advantage of statically bound systems is that errors may be detected early in the software life cycle and that checks for correct program execution may be factored out by the compiler. It is impossible to have a completely static system since we require some mechanism, e.g. the compiler, to create the bindings in the first place. Thus if the system contains that binding mechanism it already has a modicum of dynamic binding in it. This really should be obvious since to create dynamic change we need some dynamic mechanism. The modern trend in programming languages is to make as much of the binding as static as possible. The degree to which that can be achieved depends on the application.

We assert that the total spectrum of binding mechanisms is necessary for large system construction and evolution but we are unsure as to whether all of the intermediate binding mechanisms are useful. We suggest therefore that the programmer should have control over the binding mechanism.

To accommodate flexible incremental bindings we have invented the name space. This is an environment mechanism that permits the following:

- a) the storage of bindings in a name space;
- b) the dynamic use of names from a name space;
- c) the static use of names from a name space;
- d) the evolution of the names available in a name space;
- e) safe exchange of arbitrary data between parts of the system - especially between the permanent store and programs.

Name spaces are data objects which act like abstract data types to yield the values of the name space. These values are bindings between name and object and the name may be used in the current environment. To statically bind to a name space we use

with *name-space-expression* *compile statement*

which provides an environment for the statement containing all the names in the name space. The name space may be an object in the persistent store and the statement is bound in the compile time environment. The *name-space-expression* is a manifest constant, statically scoped and statically type

checked. The bindings in the name space provided for the qualified *statement* are variable or constant, manifest, statically scoped and statically type checked.

To bind a name space dynamically we use

using *name-space-expression* with *signature* do *statement*

In this the name space yielded by the *expression* must have the properties (name-type pairs) described by the *signature* in order to match correctly. It is a partial match and is the same mechanism that we use to match actual to formal type parameters in procedures and abstract data types to achieve parametric dependent polymorphism as first used in Russell[10]. In this case the *name-space-expression*, which may be regarded as an environment name, is constant but not manifest, dynamically scoped and dynamically type checked using its *signature*. The objects yielded from this name space are variable or constant, not manifest, and statically scoped and statically type checked.

We have chosen this combination of binding mechanisms to be under the programmer's control because we believe it will yield the right mix for building and maintaining large systems. For safety, most bindings will be as static as possible with some dynamic element to accommodate evolution. It is for the application designer to decide which mechanism is appropriate for a particular application.

To accommodate evolution, name spaces may grow and contract. For example

extend *name-space-expression* with *identifier-list* from *statement*

adds the new definitions of the identifier names to the name space. To contract the name space we use

drop *identifier-list* from *name-space-expression*

which will remove the identifier names from the name space.

As a name space is itself an object in the language, a name in a name space may be bound to a name space. Hence hierarchical naming structures similar to directory hierarchies are possible.

4.2 Outstanding problems

When file directories are used as an environment mechanism it is possible to write general utilities which scan a directory, often interactively, revealing the directory's contents and possibly allowing user action - e.g. delete - on each file. Some equivalent operations on name spaces may be useful, however, it is difficult to identify a primitive in terms of which they may be written.

Current practice is to step outside the type system or to depend on calling the compiler. Our desire to support more of the total programming activity within the strictly typed language leaves us confronting this problem.

5. Type Systems for Persistent Data

Ideally we would like a simple set of types, and a type algebra, so that by a succession of operations of the algebra, and provision of parameters, we could define a data type equivalent to any data model or conceptual data model. Parameterisation of such types yields schemata, i.e. the types of databases in the model, and each database is then an instance of such a type. We call this the "type alchemist's dream" as we do not yet know how to do it. In consequence the language designer has to opt for a particular model. Contrary to our original beliefs, this has to contain types motivated by long term data, because it tends to become large, and hence bulk operations and describing its regularity become important.

Based on an extensive survey[5] of existing, research and proposed languages together with experience of PS-algo[4], we have come up with the following general advice with regard to programming environments and persistent data.

- i) the data types should include a mechanism to capture user communication including graphical and audio facilities;
- ii) the data types should be expressive enough to capture the structure and regularity that will be required of the data;
- iii) the data types should include program, as procedures and abstract data types;
- iv) the data types should include a general purpose indexing mechanism, both polymorphic and variadic;
- v) the data types should include a bulk data type with operations on collections of data;
- vi) the data types should include some form of inheritance to model specialisation;
- vii) the types and bindings used should allow system evolution on an incremental and localised basis.

Our own preferences for a particular type system are defined elsewhere[6].

6. Concurrency and transactions

In the context of concurrency and transactions in the persistent information space architecture we have identified the following as outstanding problems for which solutions will need to be found. They are

- i) there should be mechanisms to permit and control shared and concurrent usage;
- ii) it should be possible to encapsulate any sequence of operations on the store into a transaction, which behaves as a single operation;
- iii) there should be privacy and access control mechanisms;

We have experimented with several versions of persistence[2,8]. Each one has a notion of transaction which allows a group of operations on data to appear atomic. The effects of the transaction are only visible to other transactions that start after the transaction has committed. Thus they are a mechanism for viewing and controlling change.

The concept of a transaction is clearly recursive with the fixed point being the machine's atomic store update. Nested transactions may be formed by grouping more than one store update into a single transaction. The semantics of these nested transactions are reasonably well understood[1] for a sequential process. At present we know of no satisfactory semantics for nested transactions in a concurrent environment.

Often associated with a transaction mechanism are methods for making the store reliable and for recovering from an aborted transaction. The stable store mechanism ensures that information is always kept (or copied) on non-volatile storage devices. Thus in the event of a system crash no information will be lost. Any stable store mechanism only gives relative protection since it may not be able to recover from some types of disaster, such as an atomic explosion. Most systems however provide an adequate level of service using a disk for the stable store. It is beyond the scope of this paper to discuss reliability except to say that our information space will be built on such a system.

Recovering from aborted transactions is much more controversial, especially in a concurrent environment. The problem is essentially one of undoing an action that other processes may have acted upon. For example having printed a document it is difficult to see how it may be unprinted. Most problems in this area involve interacting with the outside world. More about this will be said in the

next section.

For expressing concurrency there is no shortage of models. Programming languages have provided semaphores[9], conditional critical regions[13], monitors[14], path expressions[7] and the rendezvous[15] among others. There is very little agreement on which are appropriate mechanisms for all programming activity. It would appear therefore to be sensible not to build any of these into a system but to provide a method by which they may be built using a primitive and some abstraction mechanism.

7. A persistent information space architecture

Specifying a computer architecture is equivalent to specifying a programming language that defines the architecture. In recognising this we can use our rules for controlling complexity when defining the language.

Central to our aim of building a total system capable of providing for all programming activity in an integrated manner is our persistent information space. This space is made up of objects which may be simple or highly structured and are part of the universe of discourse defined by the type system of the PISA architecture. The type system must therefore be rich enough to satisfy all our requirements. This we recognise as a research challenge.

The information space is persistent. That is, the programmer has no knowledge of where the data resides. This may be done locally in main store or disk or remotely on non-local processors. The programmer is relieved of the burden of organising the physical storage of data in the system and presented with a conceptually simple model of data.


The mechanisms for binding in the persistent information space are those of the name space together with those used for introducing names in the architecture language.

At present we feel that it is premature to build in mechanisms and protocols for concurrency and transactions. We subscribe to the view that it is more sensible to build in a primitive for synchronization and a mechanism for specifying non-deterministic parallel computation and building the required protocols out of them by layers of abstraction[17]. For example it is possible to build monitors out of a primitive semaphore and a cobegin statement.

At a lower level the persistent information space is supported by a stable store mechanism for

reliability. The stable store may be distributed over many storage devices and processors and it is currently the focus of some research to build such a system.

Given the above components we envisage the architecture constructed diagrammatically as follows.



There are a number of lines built up by abstractions on the concurrency primitives. The reason for this is that although it is obvious that all of the protocols may be built out of the primitives it is not obvious that they form a hierarchy. Users of the persistent information space request a level of service for the data they are going to use. For example, two Ada tasks may request a rendezvous protocol over certain data. If this can be granted, the protocol is given to the user to manipulate the data. The type of the protocol, be it a simple lock or a complicated transaction mechanism with recovery depends on the application. The persistent information space should be regarded as being inhabited by many processes concurrently. To use data simultaneously two processes must agree on a protocol. Since the protocols contain the data any two processes that are independent in their use of

data may utilise competing protocols simultaneously.

7. Conclusions

A cause of difficulty in constructing suites of application programs out of existing languages, libraries, operating systems and databases has been identified. It may be decomposed into two parts

1. an overcomplex relationship between programs and their environment
2. inconsistency among the services provided by the environment

The approach proposed to overcome this is to design the environment and the language as a coherent whole. The initial steps outlined in this paper are sketches of the way certain facilities, such as flexible binding, reliable storage, long term storage etc. - currently provided by the operating system may be profitably specified as part of a language. Elsewhere we publish work where we demonstrate this approach is feasible.

8. Acknowledgements

Our dependence on a working implementation of PS-algol, from which we have learnt much, cannot be overstated. It was built largely by the following members of our team: Pete Bailey, Fred Brown, Paul Cockshott and Al Dearle. We also benefitted from suggestions and ideas arising in discussions with Peter Buneman and Tony Davie.

Alvey grant IKBS 104 supports this research at both Universities and at ICL. The work is also supported at Glasgow by SERC grants GRC 21977 and GRC 21960 and at St Andrews by SERC grant GRC 15907 and at both Universities by grants from ICL.

9. References

1. Atkinson, M.P., Cockshott, W.P. & Chisholm, K.J. NEPAL - The New Edinburgh Persistent Algorithmic Language. DATABASE Infotech State of the Art Report 9, 8 (1982), 299-318.
2. Atkinson, M.P., Bailey, P.J., Cockshott, W.P., Chisholm, K.J. & Morrison, R. An approach to persistent programming. Computer Journal 26, 4 (1983), 360-365.
3. Atkinson, M.P., Bailey, P.J., Cockshott, W.P. & Morrison, R. PS-algol reference manual.
- Universities of Glasgow and St Andrews PPRR-12 (1984).
4. Atkinson, M.P. & Morrison, R. Procedures as persistent data objects. ACM TOPLAS 7, 4 (1985).
5. Atkinson, M.P. & Buneman, O.P. Database programming language design. in preparation.
6. Atkinson, M.P. & Morrison, R. Types, bindings and parameters in a persistent environment. Proc. Appin workshop. Universities of Glasgow and St Andrews PPRR-16 (1985), 1-24.
7. Campbell, R.H. & Habermann, A.N. The specification of process synchronization by path expressions. Lecture Notes in Computer Science, 16 Springer-Verlag (1974).
8. Cockshott, W.P., Atkinson, M.P., Bailey, P.J., Chisholm, K.J. & Morrison, R. The persistent object management system. Software, Practice & Experience 14 (1984).
9. Dijkstra, E.W. THE multiprogramming system. Comm.ACM 11,5 (1968),341-346.
10. Demers, A. & Donahue, J. Revised report on Russell. Technical report TR79-389, (1979), Cornell University.
11. Dennis, J.R. Segmentation and the design of multiprogrammed computer systems. JACM 12,4 (1965), 589-602.
12. Gunn H.I.E.,Morrison R., On the implementation of constants Information Processing Letters 9,1 (1979), 1-4.
13. Hoare, C.A.R. Towards a theory of parallel programming. In Operating System Techniques (Ed Hoare & Perrot). Academic Press, London (1972), 61-71.
14. Hoare, C.A.R. Monitors : an operating system structuring concept. Comm.ACM 17, 10 (1974), 549-557.
15. Ichbiah et al, The Programming Language Ada Reference Manual.

ANSI/MIL-STD-1815A-1983. (1983).

16. Landin, P.J. The next 700 programming languages.
Comm.ACM 9, 3 (1966), 157-164.
17. Krablin, G.L. Building flexible multilevel transactions in a distributed persistent environment.
Proc. of the Appin workshop.
Universities of Glasgow and St Andrews PPRR-16 (1985), 83-106
18. McCarthy, J. et al. Lisp 1.5 Programmers manual.
M.I.T. Press Cambridge Mass. (1962).
19. Morrison, R. S-algol language reference manual.
University of St Andrews CS/79/1 (1979).
20. Richards M. BCPL, a tool for compiler writing and systems programming
AFIPS SJCC 1969
21. Strachey, C. Fundamental concepts in programming languages.
Oxford University Press, Oxford (1967).
22. Tennent, R.D. Language design methods based on semantic principles.
Acta Informatica 8 (1977), 97-112.
23. Turner, D.A. SASL language manual.
University of St Andrews CS/79/3 (1979).
24. Wirth, N. The programming language Pascal.
Acta Informatica 1, 1 (1971), 35-63.

Bibliography

Copies of documents in this list may be obtained by writing to:

The Secretary,
Persistent Programming Research Group,
Department of Computing Science,
University of Glasgow,
Glasgow G12 8QQ
Scotland.

Books

Davie, A.J.T. & Morrison, R.
"Recursive Descent Compiling", Ellis-Horwood Press (1981).

Atkinson, M.P. (ed.)
"Databases", Pergamon Infotech State of the Art Report, Series 9, No.8,
January 1982. (535 pages).

Cole, A.J. & Morrison, R.
"An introduction to programming with S-algol", Cambridge University Press,
Cambridge, England, 1982.

Stocker, P.M., Atkinson, M.P. & Grey, P.M.D. (eds.)
"Databases - Role and Structure", Cambridge University Press, Cambridge,
England, 1984.

Published Papers

Morrison, R.
"A method of implementing procedure entry and exit in block structured high level
languages". Software, Practice and Experience 7, 5 (July 1977), 535-537.

Morrison, R. & Podolski, Z.
"The Graffiti graphics system", Proc. of the DECUS conference, Bath (April 1978),
5-10.

Atkinson, M.P.
"A note on the application of differential files to computer aided design", ACM
SIGDA newsletter Summer 1978.

Atkinson, M.P.
"Programming Languages and Databases", Proceedings of the 4th International
Conference on Very Large Data Bases, Berlin, (Ed. S.P. Yao), IEEE, Sept. 78,
408-419. (A revised version of this is available from the University of Edinburgh
Department of Computer Science (EUCS) as CSR-26-78).

Atkinson, M.P.
"Progress in documentation: Database management systems in library
automation and information retrieval", Journal of Documentation Vol.35, No.1,
March 1979, 49-91. Available as EUCS departmental report CSR-43-79.

Gunn, H.I.E. & Morrison, R.
"On the implementation of constants", Information Processing Letters 9, 1 (July
1979), 1-4.

Atkinson, M.P.
"Data management for interactive graphics", Proceedings of the Infotech State of
the Art Conference, October 1979. Available as EUCS departmental report
CSR-51-80.

Atkinson, M.P. (ed.)
"Data design", Infotech State of the Art Report, Series 7, No.4, May 1980.

Morrison, R.
"Low cost computer graphics for micro computers", Software Practice and
Experience, 12, 1981, 767-776.

Atkinson, M.P., Chisholm, K.J. & Cockshott, W.P.
"PS-algol: An Algol with a Persistent Heap", ACM SIGPLAN Notices Vol.17, No.
7, (July 1981) 24-31. Also as EUCS Departmental Report CSR-94-81.

Atkinson, M.P., Chisholm, K.J. & Cockshott, W.P.
"Nepal - the New Edinburgh Persistent Algorithmic Language", in Database,
Pergamon Infotech State of the Art Report, Series 9, No.8, 299-318 (January
1982) - also as EUCS Departmental Report CSR-90-81.

Morrison, R.
"S-algol: a simple algol", Computer Bulletin II/31 (March 1982).

Morrison, R.
"The string as a simple data type", *Sigplan Notices*, Vol.17,3, 46-52, 1982.

Atkinson, M.P., Bailey, P.J., Chisholm, K.J., Cockshott, W.P. & Morrison, R.
"Progress with Persistent Programming", presented at CREST course UEA, September 1982, revised in "Databases - Role and Structure", see PPRR-8-84.

Morrison, R.
"Towards simpler programming languages: S-algol", *IUCC Bulletin* 4, 3 (October 1982), 130-133.

Atkinson, M.P., Chisholm, K.J. & Cockshott, W.P.
"Problems with persistent programming languages", presented at the Workshop on programming languages and database systems, University of Pennsylvania. October 1982. Circulated (revised) in the Workshop proceedings 1983, see PPRR-2-83.

Atkinson, M.P.
"Data management", in *Encyclopedia of Computer Science and Engineering* 2nd Edition, Ralston & Meek (editors) January 1983. van Nostrand Reinhold.

Atkinson, M.P., Chisholm, K.J. & Cockshott, W.P.
"Algorithms for a Persistent Heap", *Software Practice and Experience*, Vol.13, No.3, 259-272 (March 1983). Also as EUCS Departmental Report CSR-109-82.

Atkinson, M.P., Chisholm, K.J. & Cockshott, W.P.
"CMS - A chunk management system", *Software Practice and Experience*, Vol.13, No.3 (March 1983), 273-285. Also as EUCS Departmental Report CSR-110-82.

Atkinson, M.P., Bailey, P.J., Chisholm, K.J., Cockshott, W.P. & Morrison, R.
"Current progress with persistent programming", presented at the DEC workshop on Programming Languages and Databases, Boston, April 1983.

Atkinson, M.P., Bailey, P.J., Chisholm, K.J., Cockshott, W.P. & Morrison, R.
"An approach to persistent programming", *The Computer Journal*, 1983, Vol.26, No.4, 360-365 - see PPRR-2-83.

Atkinson, M.P., Bailey, P.J., Chisholm, K.J., Cockshott, W.P. & Morrison, R.
"PS-algol a language for persistent programming", 10th Australian Computer Conference, Melbourne, Sept. 1983, 70-79 - see PPRR-2-83.

Morrison, R., Weatherill, M., Podolski, Z. & Bailey, P.J.
"High level language support for 3-dimension graphics", *Eurographics Conference Zagreb*, North Holland, 7-17, Sept. 1983. (ed. P.J.W. ten Hagen).

Cockshott, W.P., Atkinson, M.P., Chisholm, K.J., Bailey, P.J. & Morrison, R.
"POMS : a persistent object management system", *Software Practice and Experience*, Vol.14, No.1, 49-71, January 1984.

Kulkarni, K.G. & Atkinson, M.P.
"Experimenting with the Functional Data Model", in *Databases - Role and Structure*, Cambridge University Press, Cambridge, England, 1984.

Atkinson, M.P. & Morrison, R.
"Persistent First Class Procedures are Enough", *Foundations of Software Technology and Theoretical Computer Science* (ed. M. Joseph & R. Shyamasundar) Lecture Notes in Computer Science 181, Springer Verlag, Berlin (1984).

Atkinson, M.P., Bocca, J.B., Elsey, T.J., Fiddian, N.J., Flower, M., Gray, P.M.D. Gray, W.A., Hepp, P.E., Johnson, R.G., Milne, W., Norrie, M.C., Omololu, A.O., Oxborrow, E.A., Shave, M.J.R., Smith, A.M., Stocker, P.M. & Walker, J.
"The Proteus distributed database system", proceedings of the third British National Conference on Databases, (ed. J. Longstaff), BCS Workshop Series, Cambridge University Press, Cambridge, England, (July 1984).

Atkinson, M.P. & Morrison, R.
"Procedures as persistent data objects", *ACM TOPLAS* 7, 4, 539-559, (Oct. 1985) - see PPRR-9-84.

Morrison, R., Bailey, P.J., Dearle, A., Brown, P. & Atkinson, M.P.
"The persistent store as an enabling technology for integrated support environments", 8th International Conference on Software Engineering, Imperial College, London (August 1985), 166-172 - see PPRR-15-85.

Atkinson, M.P. & Morrison, R.
"Types, bindings and parameters in a persistent environment", proceedings of Data Types and Persistence Workshop, Appin, August 1985, 1-24 - see PPRR-16-85.

Davie, A.J.T.
"Conditional declarations and pattern matching", proceedings of Data Types and Persistence Workshop, Appin, August 1985, 278-283 - see PPRR-16-85.

Krablin, G.L.
"Building flexible multilevel transactions in a distributed persistent environment", proceedings of Data Types and Persistence Workshop, Appin, August 1985, 86-117 - see PPRR-16-85.

Buneman, O.P.
"Data types for data base programming", proceedings of Data Types and Persistence Workshop, Appin, August 1985, 291-303 - see PPRR-16-85.

Cockshott, W.P.
"Addressing mechanisms and persistent programming", proceedings of Data Types and Persistence Workshop, Appin, August 1985, 363-383 - see PPRR-16-85.

Norrie, M.C.

"PS-algol: A user perspective", proceedings of Data Types and Persistence Workshop, Appin, August 1985, 399-410 - see PPRR-16-85.

Owoso, G.O.

"On the need for a Flexible Type System in Persistent Programming Languages", proceedings of Data Types and Persistence Workshop, Appin, August 1985, 423-438 - see PPRR-16-85.

Morrison, R., Brown, A.L., Bailey, P.J., Davie, A.J.T. & Dearle, A.

"A persistent graphics facility for the ICL PERQ", Software Practice and Experience, Vol.14, No.3, (1986) - see PPRR-10-84.

Atkinson, M.P. and Morrison R.

"Integrated Persistent Programming Systems", proceedings of the 19th Annual Hawaii International Conference on System Sciences, January 7-10, 1986 (ed. B. D. Shriver), vol IIA, Software, 842-854, Western Periodicals Co., 1300 Rayman St., North Hollywood, Calif. 91605, USA - see PPRR-19-85.

Atkinson, M.P., Morrison, R. and Pratten, G.D.

"A Persistent Information Space Architecture", proceedings of the 9th Australian Computing Science Conference, January, 1986 - see PPRR-21-85.

Kulkarni, K.G. & Atkinson, M.P.

"EFDM : Extended Functional Data Model", The Computer Journal, Vol.29, No.1, (1986) 38-45.

Atkinson, M.G., Morrison, R. & Pratten G.D.

"Designing a Persistent Information Space Architecture", proceedings of Information Processing 1986, Dublin, September 1986, (ed. H.J. Kugler), 115-119, North Holland Press.

Morrison R., Dearle, A., Brown, A. & Atkinson M.P.; "An integrated graphics programming environment", Computer Graphics Forum, Vol. 5, No. 2, June 1986, 147-157 - see PPRR-14-86.

Buneman, O.P. & Atkinson, M.P.

"Inheritance and Persistence in Database Programming Languages"; proceedings ACM SIGMOD Conference 1986, Washington, USA May 1986; The Association for Computing Machines, 11 West 42nd St., New York, NY 10036; - see PPRR-22-86.

Brown, A.L. & Dearle, A.

"Implementation Issues in Persistent Graphics"; University Computing, Vol. 8, NO. 2, (Summer 1986) - see PPRR-23-86.

Internal Reports

Morrison, R.

"S-Algol language reference manual", University of St Andrews CS-79-1, 1979.

Bailey, P.J., Maritz, P. & Morrison, R.

"The S-algol abstract machine", University of St Andrews CS-80-2, 1980.

Atkinson, M.P., Hepp, P.E., Ivanov, H., McDuff, A., Proctor, R. & Wilson, A.G.

"EDQUSE reference manual", Department of Computer Science, University of Edinburgh, September 1981.

Hepp, P.E. and Norrie, M.C.

"RAQUEL: User Manual", Department of Computer Science Report CSR-188-85, University of Edinburgh.

Norrie, M.C.

"The Edinburgh Node of the Proteus Distributed Database System", Department of Computer Science Report CSR-191-85, University of Edinburgh.

In Preparation

Kulkarni, K.G. & Atkinson, M.P.

"EFDM : A DBMS based on the functional data model", to be submitted.

Atkinson, M.P. & Buneman, O.P.

"Database programming languages design", submitted to ACM Computing Surveys - see PPRR-17-85.

Persistent Programming Research Reports

Theses

The following Ph.D. theses have been produced by member of the group and are available from:

The Secretary,
Persistent Programming Group,
University of Glasgow,
Department of Computing Science,
Glasgow G12 8QQ,
Scotland.

W.P. Cockshott
Orthogonal Persistent, University of Edinburgh, February 1983.

K.G. Kulkarni
Evaluation of Functional Data Models for Database Design and Use, University of Edinburgh, 1983.

P.E. Hepp
A DBS Architecture Supporting Coexisting Query Languages and Data Models, University of Edinburgh, 1983.

G.D.M. Ross
Virtual Files: A Framework for Experimental Design, University of Edinburgh, 1983.

G.O. Owoso
Data Description and Manipulation in Persistent Programming Languages, University of Edinburgh, 1984.

This series was started in May 1983. The following list gives those produced and those planned plus their status at 28th October 1986.

Copies of documents in this list may be obtained by writing to:

The Secretary,
The Persistent Programming Research Group,
Department of Computing Science,
University of Glasgow,
Glasgow G12 8QQ.

PPRR-1-83	The Persistent Object Management System - Atkinson,M.P., Chisholm, K.J. and Cockshott, W.P.	£1.00
PPRR-2-83	PS-algol Papers: a collection of related papers on PS-algol - Atkinson, M.P., Bailey, P., Cockshott, W.P., Chisholm, K.J. and Morrison, R.	£2.00
PPRR-4-83	The PS-algol reference manual - Atkinson, M.P., Bailey, P., Cockshott, W.P., Chisholm, K.J. and Morrison, R.	£2.00
PPRR-5-83	Experimenting with the Functional Data Model - Atkinson, M.P. and Kulkarni, K.G.	£1.00
PPRR-6-83	A DBS Architecture supporting coexisting user interfaces: Description and Examples - Hepp, P.E.	£1.00
PPRR-7-83	EFDM - User Manual - K.G.Kulkarni	£1.00
PPRR-8-84	Progress with Persistent Programming - Atkinson,M.P., Bailey, P., Cockshott, W.P., Chisholm, K.J. and Morrison, R.	£2.00
PPRR-9-84	Procedures as Persistent Data Objects - Atkinson, M.P., Bailey, P., Cockshott, W.P., Chisholm, K.J. and Morrison, R.	£1.00
PPRR-10-84	A Persistent Graphics Facility for the ICL PERQ - Morrison, R., Brown, A.L., Bailey, P.J., Davie, A.J.T. and Dearle, A.	£1.00

PPRR-11-85	PS-algol Abstract Machine Manual	£1.00		
PPRR-12-86	PS-algol Reference Manual - third edition	£2.00		
PPRR-13-85	CPOMS - A Revised Version of The Persistent Object Management System in C - Brown, A.L. and Cockshott, W.P.	£2.00		
PPRR-14-86	An Integrated Graphics Programming Environment - second edition - Morrison, R., Brown, A.L., Dearle, A. and Atkinson, M.P.	£1.00		
PPRR-15-85	The Persistent Store as an Enabling Technology for Integrated Project Support Environment - Morrison, R., Dearle, A., Bailey, P.J., Brown, A.L. and Atkinson, M.P.	£1.00		
PPRR-16-85	Proceedings of the Persistence and Data Types Workshop, Appin, August 1985 - ed. Atkinson, M.P., Buneman, O.P. and Morrison, R.	£15.00		
PPRR-17-85	Database Programming Language Design - Atkinson, M.P. and Buneman, O.P.	£3.00		In Preparation
PPRR-18-85	The Persistent Store Machine - Cockshott, W.P.	£2.00		Some Applications Programmed in a Persistent Language - Cooper, R.L. (ed).
PPRR-19-85	Integrated Persistent Programming Systems - Atkinson, M.P. and Morrison, R.	£1.00		PS-algol Applications Programming - Cooper, R.L.
PPRR-20-85	Building a Microcomputer with Associative Virtual Memory - Cockshott, W.P.	£1.00		A Compilation Technique for a Block Retention Language - Cockshott, W.P. and Davie, A.J.T.
PPRR-21-85	A Persistent Information Space Architecture - Atkinson, M.P., Morrison, R. and Pratten, G.D.	£1.00		Thoughts on Concurrency - Wai, F.
PPRR-22-86	Inheritance and Persistence in Database Programming Languages - Buneman, O.P. and Atkinson, M.P.	£1.00		Concurrency in Persistent Programming Languages - Krablin, G.K.
PPRR-23-86	Implementation Issues in Persistent Graphics - Brown, A.L. and Dearle, A.	£1.00		Providing Database Interfaces Within A Persistent Environment - Atkinson, M.P. & Cooper, R.L.
PPRR-24-86	Using a Persistent Environment to Maintain a Bibliographic Database - Cooper, R.L., Atkinson, M.P. & Blott, S.M.	£1.00		
PPRR-26-86	Exception Handling in a Persistent Programming Language - Philbrow, P & Atkinson M.P.	£1.00		
PPRR-28-86b	A Domain Theoretic Approach to Higher-Order Relations - Buneman, O.P. & Ochari, A.	£1.00		
PPRR-29-86	A Persistent Store Garbage Collector with Statistical Facilities - Campin, J. & Atkinson, M.P	£1.00		
PPRR-30-86	Data Types for Data Base Programming - Buneman, O.P.	£1.00		