

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



Building a Microcomputer with
Associative Virtual Memory

Persistent Programming
Research Report 20

Alan Dcarle .

Building a Microcomputer with Associative Virtual
Memory

W P Cockshott

Glasgow University Computer Science Dept

ABSTRACT

This report describes the motivation for and design of the POPPY computer designed to support persistent programming languages. The computer uses a novel virtual memory architecture to support object addressing.

November 14, 1985

Building a Microcomputer with Associative Virtual Memory

W P Cockshott

Glasgow University Computer Science Dept

1. Persistent Programming

Conventional programming languages like Basic, Pascal or C present the programmer with two broadly different views of memory: program variables and files. The distinction between them arose because a computer generally has two different types of physical store : RAM and Disks.

Variables are used to store the working data of a program and are held in RAM, whereas files are held on Disk and are used to transmit information from one program to another. As most programmers know to their cost, files are a lot more tedious to program with than variables. It is not just that they are slower than variables, they are also a lot less flexible.

In a language like Pascal, your variables are Typed and you can create new types of data to handle new applications. For graphics work you can define types for PICTURES, WINDOWS, COLOURS etc, whose properties are optimised to suit the algorithms that they are going to be used in. Your variables can also be structured using type constructors such as arrays, sets, records and pointers. With these facilities you can readily create arbitrarily complex data structures. The freedom that this gives the programmer is even greater in those languages that allow the dynamic creation of data on a heap. In addition, a modern programming language will provide its variables with facilities for information hiding by means of modules and scope rules.

Contrasting this with what we get from files we see that they come a very poor second. They are untyped for a start. When you open a file you do not know what its internal structure is going to be. It may have been written as a file of records, but you can open it as a file of characters, so any type rules go out of the window.

Even worse, not all types of data in files. It is easy to build up a binary tree or a linked list on the Pascal heap. If this is then written out to a file of records, then all of the pointer structures get destroyed. If you read the records back in, there is no way of linking them back together again.

To cap it all, we find that files all appear to be "global". A file created by one module of a program can be read in any other module, making it very difficult to implement information hiding between modules.

With all these disadvantages why do we put up with files?

It seems that there are three reasons. The first is just historical inertia. The distinction between files and variables has been around so long that it is part of the established paradigm of computing. It requires a conceptual revolution to break away from it. Behind this there are economic and technical reasons. Ram chips are volatile. They do not store information in the absence of a power supply, so that disk files tend to be used for information that has to persist over long periods. In addition they are more expensive than disks. Small portable computers do use battery backed CMOS Ram chips for persistent storage, but for price reasons the capacity on these is very small. However these hardware problems are by no means insuperable.

1.1. Virtual memory

Hardware designers were quick to see the possibilities inherent in rotating magnetic storage devices. Virtual memory using paging techniques was achieved by the Manchester University "Atlas" computer in 1960. This allows rotating storage devices to be seen as an extension to the (Von Neumann) random access store.

In the quarter century since that basic advance, software designers have made surprisingly little use of the potentialities of virtual memory. It has been seen mainly as a way of being able to run bigger programs. An area of disk is set aside for page swapping, but this inherits both the advantages and the disadvantages of conventional RAM. It is viewed as both random access, and volatile. Persistent data is still kept in files.

1.2. Store mapping

Some of the relatively early virtual memory operating systems, like Multics and the lesser known Edinburgh Multi Access System provided an advance, in that they allowed files to be mapped into the random access virtual memory. Persistent data could then be accessed as if it were an array by the normal operations of a programming language. However, this approach was only supported in certain systems programming languages and never came into wide use. The operating systems still maintained two separate types of store, governed by different conventions. There was the file store which was public, hierarchically organised and addressed associatively via long symbolic file names. Then there was the virtual memory which was essentially private and organised as an array. One area of store could be partially mapped onto the other, but that presupposed that they were seen as different in the first place.

1.3. PS-algol

From the late '70s there have been breakthroughs in programming language design that enable us at last to get away from the old file/variable paradigm. This has come through the incorporation of the idea of persistence into a number of experimental programming

languages. The first of these were Smalltalk and PS-algol, former developed at Xerox PARC and the latter at Edinburgh and St Andrews universities. Subsequently the idea has been adopted in a number of other languages: B from the Mathematical Centre in Amsterdam, Poly from Cambridge University, and Amber from Bell Labs.

These languages allow variables of any type to persist and have no notion of files in the usual sense.

PS-algol for example supports a persistent heap. Data of any type including procedures can be put on the heap, which will then persist beyond the time of execution of the program that created the data. The important thing that distinguishes a persistent heap is that pointers or references to objects can be made to persist, so that a linked list or tree used by one program can be accessed weeks later by another. The PS-algol system itself hides the existence of different store models in the underlying operating system but only at the cost of considerable software complexity and associated performance overheads. The various teams implementing Smalltalk have also found that although a persistent heap provides a superb programming environment, the complex software needed to sustain it means that you need very fast processors to get an acceptable performance.

1.4. Virtual Memory Micros

Over the last couple of years however, it has become possible to buy microprocessors with virtual memory. It should thus be possible to build a relatively cheap personal computer with hardware support for persistent programming. The rest of this article describes the Poppy, a single board computer with a virtual memory system optimised for persistent programming languages.

What are the requirements of persistent programming that have made it relatively difficult to support with purely software techniques?

Persistent programming systems, whether PS-algol or Smalltalk, are object based. Instead of viewing store as a uniform array of bytes, they view it as a collection of objects on a heap. Each object is addressed via a unique Persistent Identifier (PID). A PID is valid not just for the run of one program like a normal store address, but for an indefinite period. Objects identified in this way may migrate through the store layers from RAM to disk and back again.

This means that there are potentially a very large number of PIDs and secondly that since there is no fixed relationship between a PID and a store location, the underlying addressing mechanism is associative rather than direct.

BOX 0

Object Oriented Programming

In object oriented languages like PS-algol or Smalltalk objects may be simple things like strings and arrays or they may be more active objects capable of carrying out computation in their own right. An example of this might be an abstract object like a dictionary that maps words in one language into those of another. The following program fragment shows how we might define this sort of object in PS-algol.

1) First we must say what interface the dictionary object will present to the outside world. For this PS-algol uses what it calls "Structure Classes".

```
structure dictionary(
  proc(string->string) trans;
  proc(string,string) define
)
```

This defines a dictionary to be an object that has two attributes, an ability to map strings into strings and to define mappings between strings.

2) We now define a Generator Function that will manufacture dictionaries

```
let make.dict = proc(->pntr)
begin
  structure cell(string name, equiv:pntr left,right)
  let head:=nil

  let enter = proc( cpntr curr.cell,temp -> pntr )
  case true of
  curr.cell = nil : temp
  temp( name ) < curr.cell( name ) : begin
    curr.cell( left ) := enter( curr.cell( left ),temp )
    curr.cell
  end
  temp( name ) > curr.cell( name ) : begin
    curr.cell( right ) := enter( curr.cell( right ),temp )
    curr.cell
  end
  default : curr.cell
  end
  dictionary()
  proc(cstring n ->string)
  begin
    let p := head
    while p ~= nil and n ~= p( name ) do
      p := if n < p( name ) then p( left ) else p( right )
      if p is cell then p( equiv ) else "unknown"
    end,
    proc(string s1,s2)
    head:=enter(s1.cell(s1,s2,nil,nil))
  )
end
```

BOX 0 (continued)

This defines make.dict to be a function that creates a dictionary with the functional attributes required to define and translate words.

3) We now declare a dictionary in the global environment:

```
let english.to.american = make.dict()
```

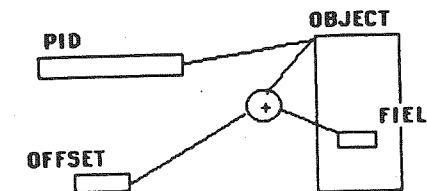
The english.to.american dictionary will then last as long as the global environment lasts, be that seconds or months.

4) We can then insert and lookup words:

```
english.to.american(define)(“colour”, “color”)
english.to.american(define)(“programme”, “program”)
write english.to.american(trans)(“dived”)
```

Box 1

Each object in the PS-algol system has a unique identifier (called a PID or Persistent IDentifier in PS-algol) and addressing has to be done in two parts, using an object ID, and an offset into an object.



1.5. Persistent Address Spaces

How big is a big address space?

One is accustomed to think of a 24 bit address space as large and a 32 bit address space as huge. Once you go over to persistent programming, things look a bit different. For a start you have to divide the address space up between different users. You could give every user their own 32 bit address space, but this would mean that they could not share data, which would not be very satisfactory. On a large machine you can easily have one or two hundred registered users. If you reserved a byte of the address to specify a user number you find that you are down to 16 megabytes per user. Remember that this has to hold not just transitory data that is used during computations, but all the long term data that would normally go into the file store.

For most users at present 16 megabytes would be enough, but for some it is already a bit tight and we know that with the development of computing the amount of store used has risen inexorably. If you start

holding digitised images or digitised voice online then you are going to need a lot more memory. A 32 bit address space is likely to look rather tight for persistent programming within the the next decade or so. But just how big an addressing system does persistent programming demand?

In languages that use a heap, a great number of objects are created that last only a short while. If each object were given a new object number, object numbers will be used up faster than the accumulation of persistent objects would justify. Broadly there are two solutions to this, either you make the object space so big that you never run out, or you build in a garbage collector that is able to recover unused object numbers.

The first solution seems to be the better. Once you go to billions of objects, conventional garbage collection techniques are likely to be too expensive. A pointer following garbage collector would be far too slow when working on a heap of this size. A reference count one would impose an overhead cost on every pointer assignment and upon every stack retraction. It is likely that on very large heaps it will be necessary to use incremental garbage collections based upon local information to tidy up local regions of the global heap. Such techniques will hold onto some things that could be discarded, but provided that we are not going to run out of object numbers, this would be supportable. Infrequently used objects would migrate onto optical archive store. Some objects in archive store will never be reused because they are actually unreachable though the system does not know it, but it is in the nature of archive stores to hold a lot junk that nobody is going to ever look at again.

How big does your object address space have to be so as not to run out of object numbers?

Back of an envelope calculations indicate that a 48 bit object number is likely to be enough. A computer creating half a million objects per second would take over 10 years to run out of object numbers, by which point it is likely to have been scrapped. In the end, in order to allow room to expand the address space accross a network, I chose a 64 bit PID (see box 2).

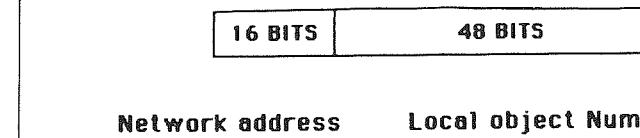
1.6. What chip to choose

There are now several microprocessors supporting virtual memory: the Intel 286 used in the IBM PC/AT, the Motorola 68020, the Intel 432 and the National Semiconductor 32000 series. Which one is most suitable for an associative, object oriented virtual memory?

Examination of all of these architectures, led me to the conclusion that only the NatSemi 32000 series met all of the requirements. At first this may seem paradoxical, as the 32000 series (like the 68000 series) has

Box 2

An obvious extension to the idea of a large persistent address space is to allow the address space to be distributed over several machines. This requires a further extension of the address space. If we allow for a local area network with 64000 machines, we find that a 64 bit key would uniquely identify every object ever created on any machine on the network.



the classic Von Neumann view of memory as a uniform array of words, whereas the Intel machines are explicitly object oriented. Both of them have serious limitations. The 286 allows far too few objects: each user only has access to 8192 object descriptors, this comes from it being basically a 16 bit machine. The intel 432 is a more serious proposition. Again it was rejected as having too small an address space. Other disadvantages of the 432 were its low speed and the lack of affordable compilers for it.

The National Semiconductor 32016 turned out to be flexible enough for some quite simple additional hardware to turn its flat, paged, virtual memory into an associative object oriented virtual memory.

1.7. Programmers view of Poppy

To the machine level programmer, the Poppy's store is divided into 3 parts:

- 1 Machine registers make up the fastest store, it has 8 floating point registers, 8 general purpose registers and 8 special purpose registers.
- 2 Temporary values that will not fit into registers are stored on a stack.
- 3 Data of longer duration is held on a vast heap, made up of a practically inexhaustible number of named objects.

The Poppy instruction set is an extended version of the NS32000 architecture. Certain of the machines addressing modes to enable them to dereference 64 bit object keys rather than 24 bit virtual addresses.

The basic NS32000 architecture supports a 24 bit virtual address space, which was well described in the April 1983 edition of Byte. Most addressing modes on the Poppy still operate using this space.

- 1 Register

The general purpose and floating point registers serve as the operands in Register address mode.

- 2 Register Relative

A general purpose register contains a 24 bit virtual address, to which a displacement is added to derive the effective address.

3 Memory Space

This is similar to the Register Relative address mode but uses one of the dedicated high level language registers PC,SP,FPSB plus an offset to get to the operand. These register again all contain 24 bit pointers.

4 Object Dereference

This is the new addressing mode added by the PSm to the basic NS32000. It allows a 64 bit PID in memory to be dereferenced. An offset is added to the start of the object referred to by the pointer to obtain the operand, it is explained in box 3.

5 Immediate

The operand is encoded within the instruction.

6 Absolute

An absolute address within the 24 bit virtual address space is specified.

7 Top of stack

The operand is Pushed (Popped) onto (from) the top of the stack.

8 Indexing

Any of the address modes other than the Register or Immediate may be further qualified by an index. Indexing has the effect of calculating an "effective address" (in virtual memory) and then multiplying one of the general purpose registers by 1,2,4,or 8 and then adding it to the total to get the final Effective address of the operand. The indexing option thus allows for arrays of bytes, short integers, integers, and PIDs or reals.

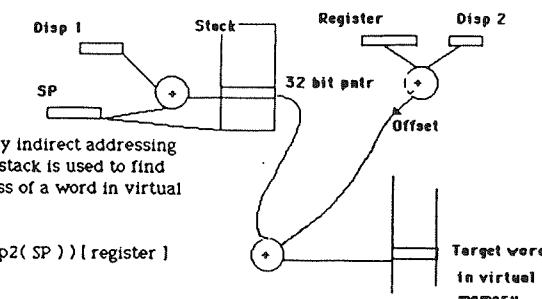
1.8. Poppy Hardware architecture

The Poppy is a single board virtual memory computer based upon the National Semiconductor 32000 series chipset, with a custom additional memory management unit. The overall logical structure of the processor is shown in the following diagram:

1.8.1. Processor

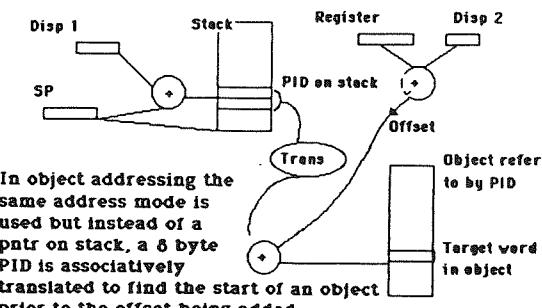
The Processor is made up of 4 National Semiconductor chips, a timer, a CPU, a memory management unit or MMU and a floating point unit. These are connected via a 16 bit multiplexed private address/data bus. In normal operation the CPU outputs a virtual address onto the

Box 3



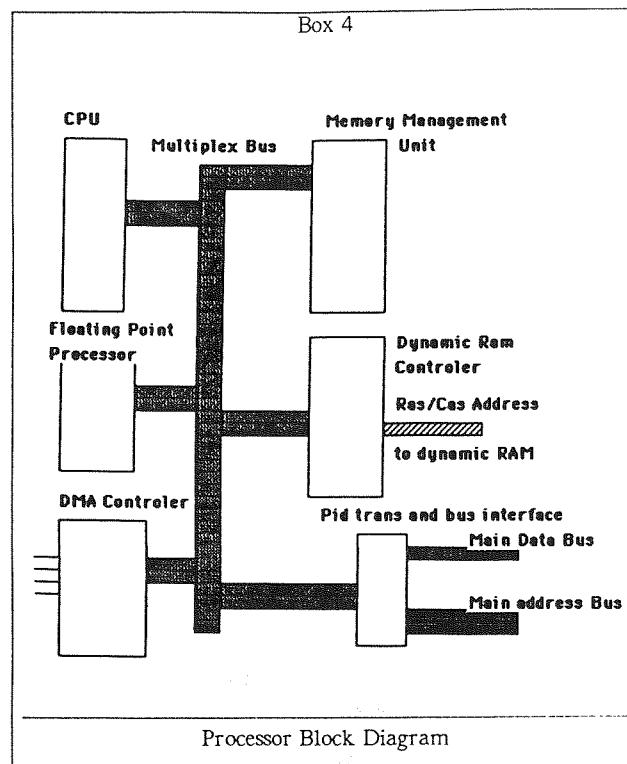
In memory indirect addressing
a pptr on stack is used to find
the address of a word in virtual
memory.

Notation:
 $disp1(disp2(SP)) [register]$



In object addressing the
same address mode is
used but instead of a
pptr on stack, a 6 byte
PID is associatively
translated to find the start of an object
prior to the offset being added

Diagram on object addressing mode



multiplexed bus during the first time interval of a store cycle. This is latched into the MMU which translates it into a physical address which is output on the next store cycle.

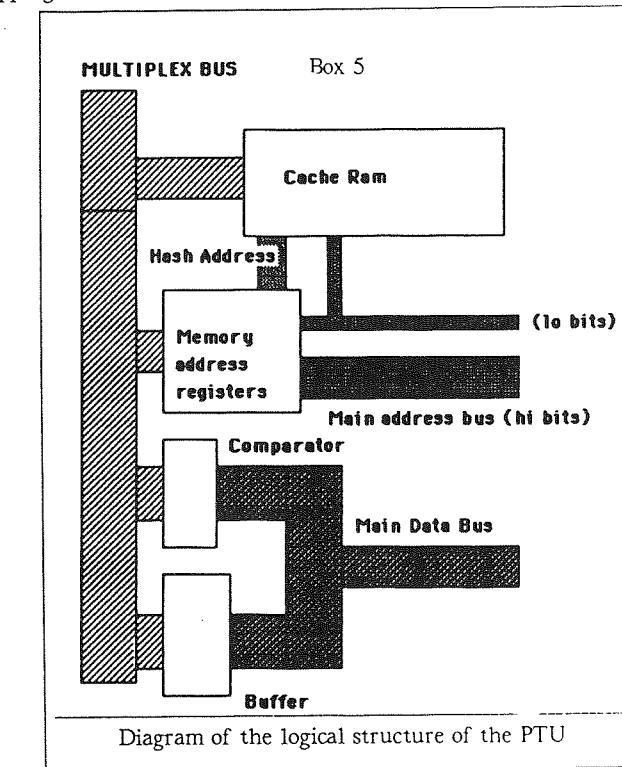
1.8.2. Memory

The physical address is latched and output to the demultiplexed address bus to which the memory is connected. This memory is made up of both nonvolatile RAM and more conventional volatile Dynamic RAM. The nonvolatile RAM is made of high speed CMOS static ram with lithium battery backup to provide 10 year data retention. The dynamic ram can be made up of either 64k or 256k chips. To save board area the dynamic RAM is mounted on 256k or 1 Meg SIPs, enabling either 512k bytes or 2 Mega bytes to fit into an area of 4.8" by 2.2". This way of mounting chips looks like becoming a standard for dynamic rams since it gives about 4 times the density of conventional DILs. Data coming from or going to the memory travels along the demultiplexed data

bus.

1.8.3. PID Translate Unit

Between the multiplexed address/data bus and the demultiplexed buses is the PID translate unit (PTU). This unit supports the new object addressing mode. Logically it is a 512 word associative cache memory mapping PIDs to virtual addresses.

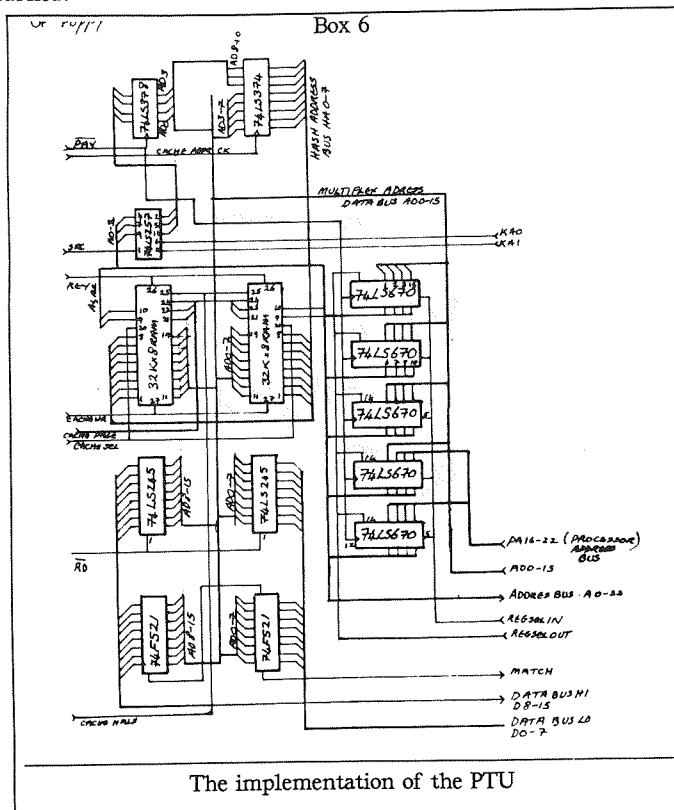


When the Processor uses the Object Indirect mode of addressing, the CPU thinks that it is just trying to fetch a 32 bit pointer from memory as per the NS32000 memory relative addressing mode. Instead of allowing this to happen, the PTU traps this, and fetches not a 32 bit pointer but a 64 bit PID from memory. It looks this up in the associative cache and returns the virtual address corresponding to the PID. The processor "thinks" it is using the old memory relative addressing mode, when in fact there is an associative operation going on behind the scenes.

How does this associative memory work?

A true associative memory chip is a set of entries each made up of M bits of tag and N bits of data. When the chip is presented with an M bit pattern on its inputs, it simultaneously compares this with the tag fields of all the entries and if one of them matches, the corresponding N bits of data are returned. Because of technical difficulties in producing such chips they are not readily available, and practical associative memories in applications like Mini or Mainframe cache memories have to be built out of conventional RAM chips.

These work on a modification of the hashing techniques used for software table lookup applications. Two banks of fast static RAM are used, the tag RAM which is M bits wide and the N bit wide data RAM. The input pattern is hashed to provide an address in the tag RAM and then the tag at that address is compared with the input tag, if they match then the data at the corresponding address in the data ram is returned.



This is the technique used in the Poppy, except that for reasons of economy the tags and the data are held on the same chips. The PTU architecture is shown in Box 5. It sits between the processor's private multiplexed address/data bus and the main address and data busses going to memory. It is made up of:

Buffer

Between the databus and the private address/data bus are a 16 bit bidirectional buffer and a 16 bit comparator connected in parallel. This buffer can be opened to allow the passage of data to and from the private bus. The comparator detects whether the contents of the two buses is the same.

The buffer is implemented with 74ls245s and the comparator with a pair of 74f521s.

Memory-Address-Registers

The private bus is connected to the address bus via two memory address registers, the Data Address Register and the Pid Address Register. These are implemented by using 2 of the 4 registers in a set of 74ls670s.

Cache

The cache is made up of 8192 by 16 bit words of CMOS static RAM.

Cache-Address-Register

The low order 3 address bit of the cache come from the main address bus, but the mid order bits come from a separate 8 bit Hash register, and the high order bits come from the control unit.

Control-Unit

This unit monitors the CPU bus transactions and steers data between the various busses and registers.

During a normal memory fetch the address is latched into the Data Address Register which is then output onto the address bus. During an object indirect instruction, the following sequence of events takes place:

- 1 The CPU outputs the virtual address of a PID on the stack, simultaneously the Status Monitoring unit spots that this is an Object Indirect addressing mode and puts the
- 2 The MMU translates this to a physical address
- 3 The PTU senses that this is a PID fetch and takes hold of the bus and latches the physical address of the PID into the Pid Address Register.
- 4 The memory returns the first 16 bits of the PID, which pass through the bus buffers and a hashing function is performed to

produce an 8 bit result. This is latched into the Hash Address Register.

- 5 The bus buffers are disabled and the PID cache outputs the first 16 bits of the PID found in the cache onto the private bus.
- 6 This is compared with the corresponding 16 bits of the PID being read from main memory.
- 7 If they are equal the next 16 bits of the PID in memory and the PID in the cache are compared, and so on until all 64 bits of the PID from memory are shown to correspond to the pid in the cache.
- 8 If the two PIDs differed, then an address fault interrupt is generated, otherwise the cache returns to the CPU the virtual address at which the object corresponding to the PID is located. The processor then adds the contents of an index register to the virtual address to find the field within object that the instruction wants.

1.8.4. Input Output

The board has two IO interfaces. There are iSBX bus sockets and there is the BBC tube. Both of these are proprietary interfaces. The two of them are under the control of DMA and interrupt controller chips, the NSC 16203 and NSC 16202 respectively. The NSC 16203 provides 4 DMA channels and the NSC 16202 provides either 8 or 16 interrupt channels.

The iSBX bus is a simple IO bus developed by Intel for their multibus series of boards. It allows small daughter boards to be mounted on a CPU board. Each of these boards may contain one or two IO devices. Several manufacturers produce these boards and the interface was chosen for its simplicity and the cheapness of the boards. Examples of the sort of functions that are available on iSBX boards are serial interfaces, parallel interfaces, floppy disk controllers, SASI interfaces and graphics boards.

There are two iSBX bus sockets on the board. Each looks like a set of 16 memory mapped byte wide registers. Each socket also has associated with it a DMA channel and an interrupt channel.

The Tube is a proprietary high speed IO interface using custom VLSI chips developed for the BBC microcomputer, manufactured by Acorn Computers. Along with carefully designed software protocols it allows the Poppy to delegate all terminal, network, and disk IO to the BBC microcomputer. High level commands are passed to the BBC machine which then executes the IO operation concurrently with the Poppy, allowing the Poppy to continue with computation. The tube is connected to a DMA and an interrupt channel on the Poppy.

The BBC micro can control floppy disks and winchester disks and provides a bit mapped graphics display, with the option of a mouse interface.

1.8.5. Database Assist Hardware

Two additional pieces of circuitry have been included to speed up database searches on the persistent store. One of these is the PF474 string search processor described in the November 1984 edition of Byte. The other is a special hashing unit that is intended to be used for computing the indices of relational databases in persistent store.

1.9. Storage management

Storage management software on the Poppy has to cooperate with the hardware in presenting the programmer with a view of single, large, object addressed store within which all distinctions between different storage media and the geographical location of these media are effaced. The storage system can be viewed from 3 levels:

Object store
Paged virtual memory
Physical store made up of:
Non volatile RAM
Volatile RAM
Rotating store

1.9.1. Object store

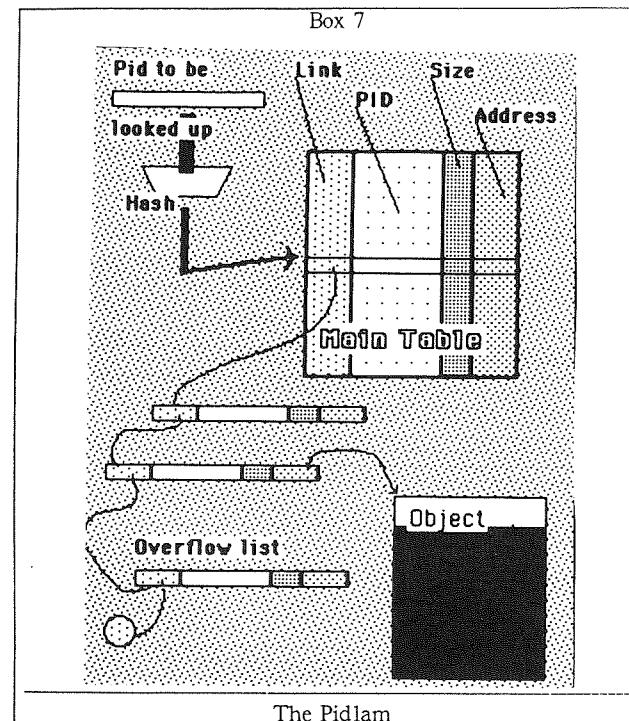
The object store is made up of up to 2 to the power of 64 distinct objects. In principle each of these should be able to contain from 1 to 4 gigabytes of data. However for reasons to do with the restriction of the addressing of early models of the National Semiconductor 16032 series, the practical limit for the size of each object is somewhat under 8 megabytes.

Each object is mapped onto the paged virtual memory of the processor as it is used, in a way analogous to the mapping of virtual pages onto physical pages in a conventional virtual memory system. On occasions objects have to be unmapped from the virtual memory of the 16032 in order to prevent the virtual memory becoming too full.

At run time the PTU will translate PIDs to virtual addresses provided that the appropriate PID/address pairs are loaded into the cache. If the PID is not found in the cache a PID/missing interrupt is generated and a software procedure is invoked to locate the object referred to by the PID and load the PID/address pair into the cache. The PID which caused the address fault has had its physical address stored in the PID Address Register. Indirection on this register enables the value of the PID itself to be located.

1.9.2. Pidlam

This software procedure uses a data structure termed the PIDLAM which is short for Persistent IDentifier to Local Address Map. The PIDLAM is a hash table in virtual memory with the structure shown in Box 7. Entries are found by Hashing a PID and then if necessary,



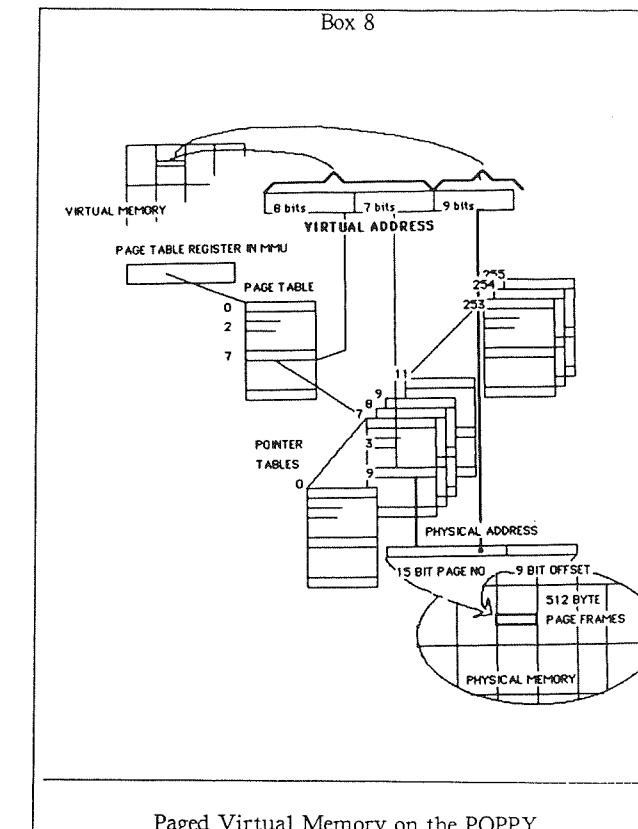
chasing down an overflow chain until an entry with the same pid is found.

There is an entry in the PIDLAM for each object currently mapped onto virtual memory. Therefore, if an object is resident in virtual memory, the PID/missing interrupt can be met by searching the PIDLAM. Otherwise the object must be taken from rotating store and mapped or moved into virtual memory.

1.9.3. Paged Store

The paged virtual memory is accessed via two level page tables as shown in Box 8. The 24 bit virtual address of the NS16032 provides 16 megabytes of virtual address space.

For reasons of efficiency, we divide objects into two great classes - the paged and the non-paged. Large objects are paged, small objects are not. This distinction arises from a desire to make the best wish of the two types of virtual memory technology on the POPPY. We have paging hardware and object addressing hardware.



The object addressing hardware maps objects to virtual addresses, the paging hardware maps virtual pages to physical pages. The simplest way to use these would be to allocate a range of virtual addresses for a heap, and map a working set of objects onto this heap. When an object was addressed and found not to be present, it would be copied from disk into the heap. This is what happens with existing software implementations of PS-algol.

This approach has a the drawbacks that if we are dealing with very large objects then we may be faced with the overhead of bringing in a lot more data than we actual need. If we just alter one word of a 10000 element array, the whole array is still copied in to the heap.

To overcome this, we chose to COPY small objects onto the heap, but to MAP large objects onto to virtual memory.

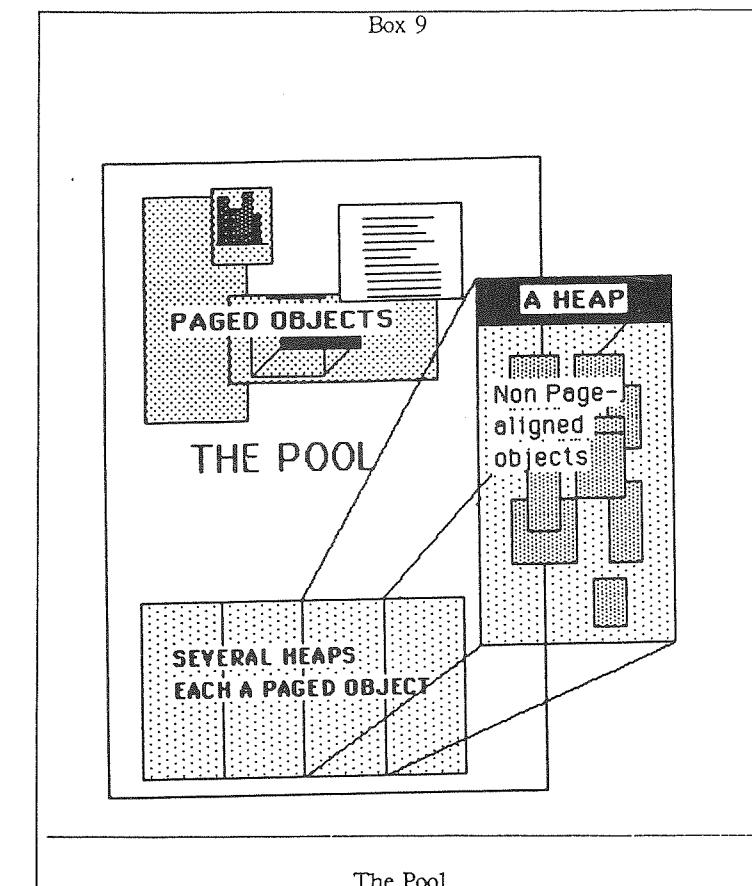
A mapped object need not all be resident in physical memory, instead it occupies a range of virtual pages individual members of which are brought into physical memory on demand. Non-paged objects are copied into a heap on demand. To the extent that the heap into which they are copied is itself paged, then they also need not be physically resident. However, it does seem reasonable to keep the virtual size of the heap sufficiently small as to ensure that most of the heap is likely to remain in physical RAM. Otherwise, we would be faced with having two swapping mechanisms competing with one another.

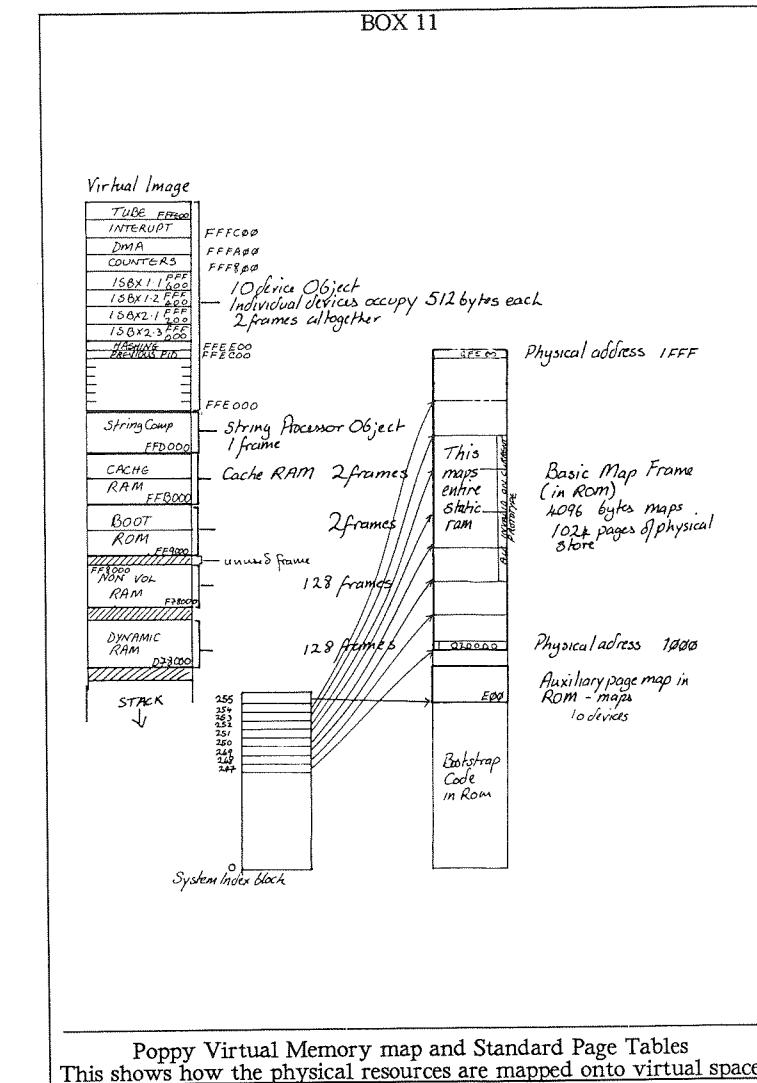
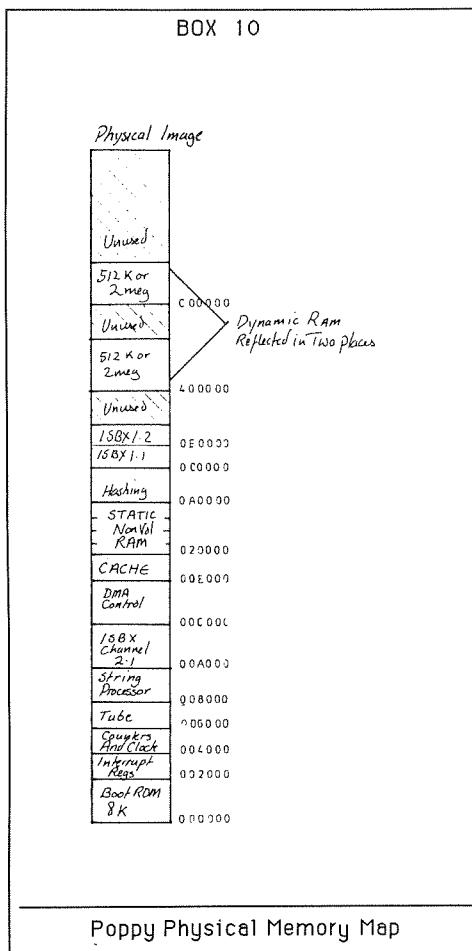
Only one copy of each object is ever present so that all transactions that can access the object work on the same copy. The definition of the location of an object in virtual memory is handled by the PIDLAM (Persistent ID to Local Address Map).

1.10. Making Object Oriented Languages really run

At least initially, the Poppy will be a single language machine, the software will consist of an interactive PS-algol compiler and its run time support. Any data declared at the outermost level of the system, whatever its type will persist indefinitely.

For some time PS-algol has been running on the ICL/Three Rivers Perq computer, which is a workstation in the same general class as the Xerox star. Because this lacks the necessary associative addressing hardware, the language runs slowly. The experience of Smalltalk implementations too is that it is difficult to get a satisfactory performance unless you have a very high powered machine like the Dorado (which costs tens of thousands of dollars). Hopefully, the type of simple associative hardware used on the Poppy should enable these sophisticated languages to be run fast on the next generation of cheap personal computers.





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.

Atkinson, M.P., Chisholm, K.J. & Cockshott, W.P.
"Towards simpler programming languages: S-algol", ICCC 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.

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 : Extended Functional Data Model", to be published in The Computer Journal.

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.

Morrison, R., Dearle, A., Bailey, P., Brown, A. & Atkinson, M.P.

"An integrated graphics programming system", to be presented at EUROGRAPHICS UK, Glasgow University, March 1986 - see PPRR-14-86.

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.

Persistent Programming Research Reports

This series was started in May 1983. The following list gives those produced and those planned plus their status at 17th March 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.	[Printed] £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.	[Printed] £2.00
PPRR-3-83	The PS-algol implementor's guide	[Withdrawn]
PPRR-4-83	The PS-algol reference manual - Atkinson, M.P., Bailey, P., Cockshott, W.P., Chisholm, K.J. and Morrison, R.	[Printed] £2.00
PPRR-5-83	Experimenting with the Functional Data Model - Atkinson, M.P. and Kulkarni, K.G.	[Printed] £1.00
PPRR-6-83	A DBS Architecture supporting coexisting user interfaces: Description and Examples - Hepp, P.E.	[Printed] £1.00
PPRR-7-83	EFDM - User Manual - K.G.Kulkarni	[Printed] £1.00
PPRR-8-84	Progress with Persistent Programming - Atkinson, M.P., Bailey, P., Cockshott, W.P., Chisholm, K.J. and Morrison, R.	[Printed] £2.00
PPRR-9-84	Procedures as Persistent Data Objects - Atkinson, M.P., Bailey, P., Cockshott, W.P., Chisholm, K.J. and Morrison, R.	[Printed] £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.	[Printed] £1.00

PPRR-11-85	PS-algol Abstract Machine Manual	[Printed] £1.00
PPRR-12-85	PS-algol Reference Manual - second edition	[Printed] £2.00
PPRR-13-85	CPOMS - A Revised Version of The Persistent Object Management System in C - Brown, A.L. and Cockshott, W.P.	[Printed] £2.00
PPRR-14-86	An Integrated Graphics Programming Environment - second edition - Morrison, R., Brown, A.L., Dearle, A. and Atkinson, M.P.	[Printed] £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.	[Printed] £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.	[Printed] £15.00
PPRR-17-85	Database Programming Language Design - Atkinson, M.P. and Buneman, O.P.	[Printed] £3.00
PPRR-18-85	The Persistent Store Machine - Cockshott, W.P.	[Printed] £2.00
PPRR-19-85	Integrated Persistent Programming Systems - Atkinson, M.P. and Morrison, R.	[Printed] £1.00
PPRR-20-85	Building a Microcomputer with Associative Virtual Memory - Cockshott, W.P.	[Printed] £1.00
PPRR-21-85	A Persistent Information Space Architecture - Atkinson, M.P., Morrison, R. and Pratten, G.D.	[Printed] £1.00
PPRR-22-86	Some Applications Programmed in a Persistent Language - Cooper, R.L., Cranston, R.D., Dearle, A. and MacFarlane, D.K.	[In Preparation]
PPRR-23-86	PS-algol Applications Programming - Cooper, R.L., Dearle, A., MacFarlane, D.K. and Philbrow, P.	[In Preparation]
PPRR-24-86	A Compilation Technique for a Block Retention Language - Cockshott, W.P. and Davie, A.J.T.	[In Preparation]
PPRR-25-86	Thoughts on Concurrency - Wai, F.	[In Preparation]

