

THE EVOLUTION OF THE OPERATING SYSTEM EMAS 2900

P.D. Stephens and J.K. Yarwood

**Edinburgh Regional Computing Centre,
Mayfield Road, Edinburgh 9**

D.J. Rees and N.H. Shelness

**Department of Computer Science,
University of Edinburgh,
Mayfield Road, Edinburgh 9**

SUMMARY

As a result of experiencing problems with manufacturer's software on the early 2900 machines Edinburgh University took the bold step of attempting to move the general purpose time-sharing system EMAS from an ICL 4-75 to an ICL 2970 computer. This paper describes the move together with the changes made and includes some preliminary performance figures from the new system.

This implementation project has produced a major time-sharing virtual memory operating system with a fraction of the effort required to implement the original system.

Keywords: 2900 Series, Operating System, Multi-Access, Benchmarking, Re-implementation.

INTRODUCTION

By January 1974 the Edinburgh Multi Access System (EMAS)¹³ had reached a high level of efficiency. When running on the ICL 4-75 machine whose configuration is shown in Figure 1 it was capable of supporting up to 32 simultaneous terminals with satisfactory response. The system was popular and soon overloaded and the Regional Computing Organisation* drew up an Operational Requirement for a further machine; this machine was to support 64 terminals (i.e. about twice those supported on the 4-75) and simultaneously process twice the batch throughput of the University 370/155.

* The Regional Computing Organisation comprises the Universities of Edinburgh, Glasgow and Strathclyde.




Figure 1 EMAS 4-75 Configuration

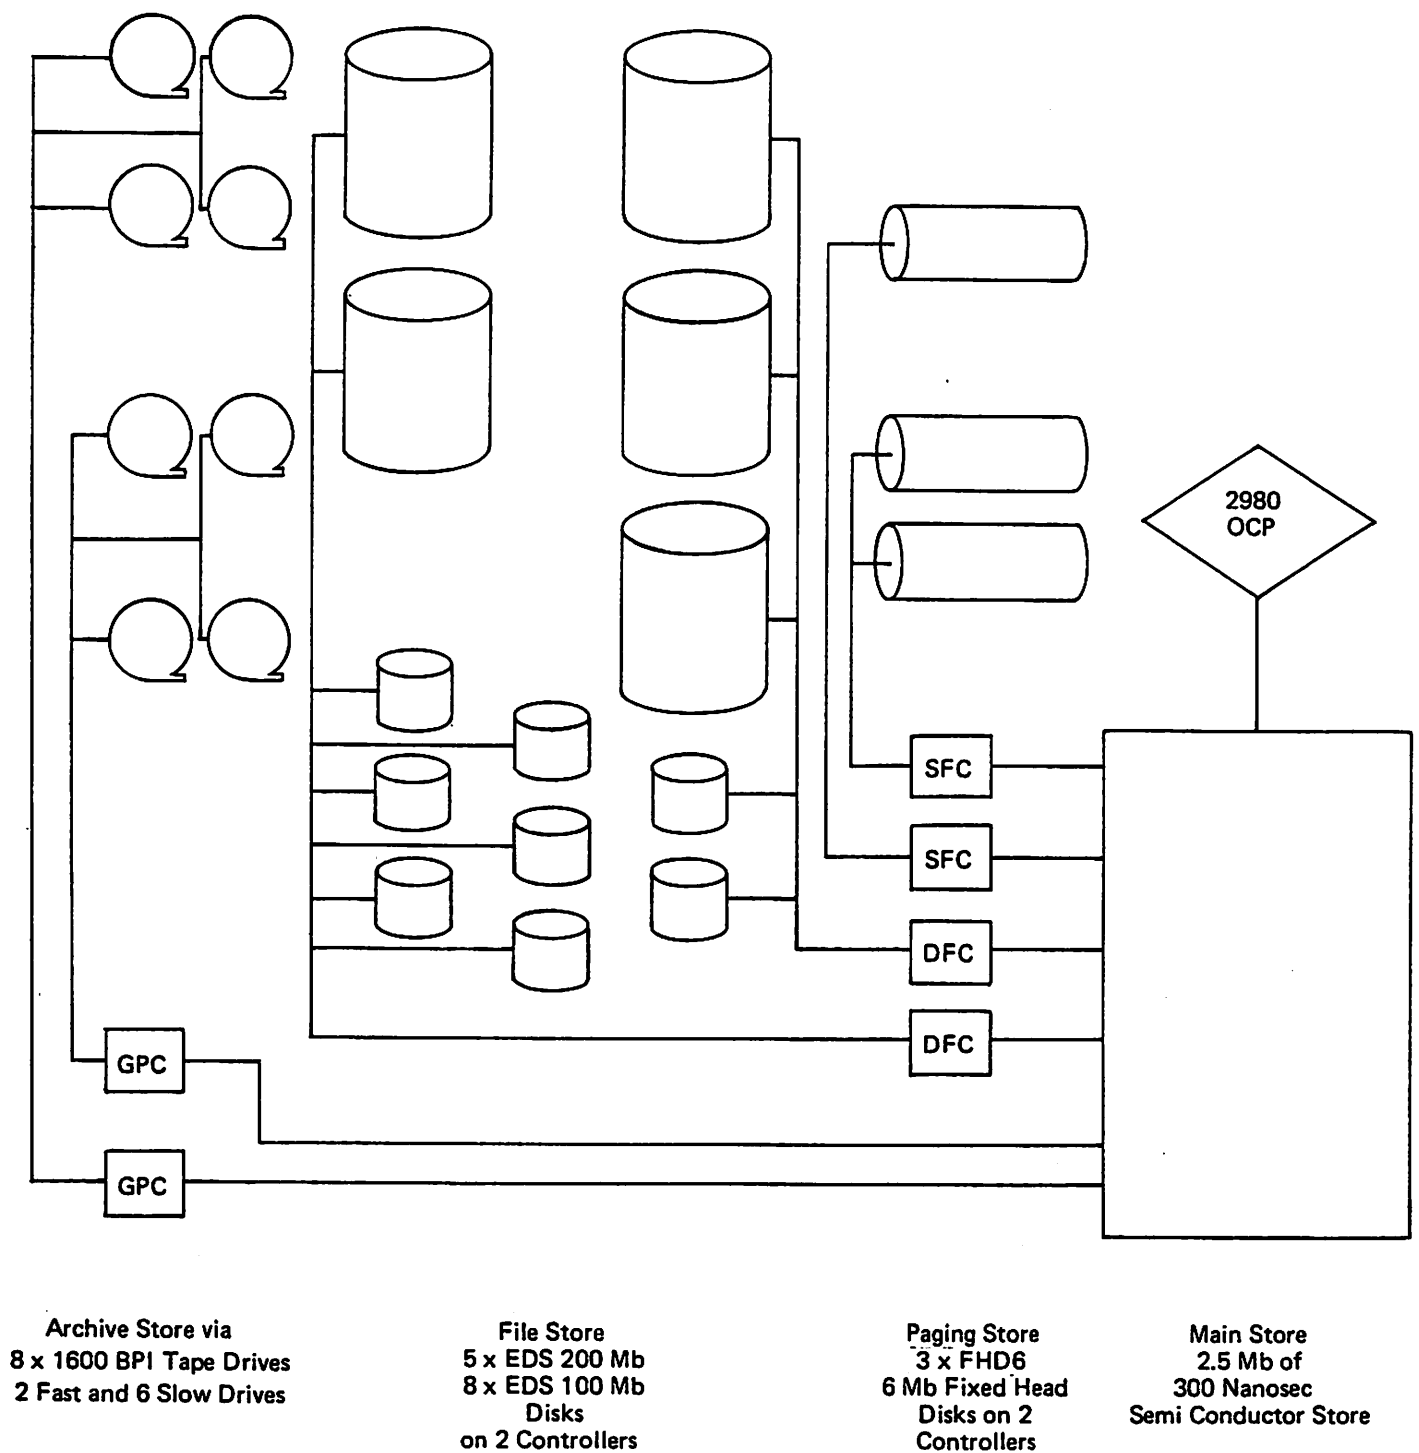



Figure 2 EMAS 2980 Configuration



Archive Store via
4 x 1600 BPI Tape Drives

File Store
4 x EDS 100 Mb
Disks
on 1 Controller

Paging Store
2 x FHD6
6 Mb Fixed-head
disks on 1 Controller

Main Store
1 Mb of
600 Nanosec
Semi-Conductor
Store

Figure 3 EMAS 2970 Configuration

A four-part batch and interactive benchmark (the Glasgow Benchmark)¹¹ was drawn up based partly on performance measurement of EMAS.

The RCO accepted a tender from ICL for delivery of an ICL 2980 in September 1975 with an undertaking to pass all four parts of the agreed benchmark by June 1977. The configuration of the 2980 is shown in Figure 2. In fact the 2980 destined for the RCO was diverted to the European Space Agency and a 2980 did not arrive until May 1976. A smaller machine - the ICL 2970 (configuration in Figure 3) - arrived in early summer 1975 for evaluation and software development. At this stage there was no intention of writing an operating system and no relevant work was done apart from writing a compiler for the EMAS implementation language IMP¹². By early summer 1976 most of Edinburgh's system programmers were convinced that ICL could not pass the interactive benchmark by a large margin and began to speculate idly on the possibilities of bootstrapping an operating system from a 4-75 to a 2970. ICL remained confident, the University authorities appeared to believe ICL and inaction followed. It was not until October '76 that a serious attempt was made to produce an alternative system. A small team of four full-time people and about half a dozen part-time helpers set out to write the EMAS 2900* operating system. Progress was rapid. By Spring '78 an internal service for project members was working, and a few "real" users were invited to use it. In October '78 a service was offered on the 2970 although development continued. In January 1980 service was opened on the RCO 2980 and also on the University of Kent 2960.

* In this paper EMAS refers to the original system running on the 4-75 and EMAS 2900 to the re-implemented system.

We do not propose in this paper to give a detailed introduction to the design or initial implementation of EMAS. An overview has been published¹³, together with detailed descriptions of specific aspects of the system^{2,7,8,10}. It is useful to note the features that give the system its current appearance:

1) Interactive working

The system is normally accessed from an interactive terminal, which it uses as its major source of control information. This may be one of a large number of devices attached via a geographically distributed network.

2) Multiple virtual memories

Each user is allocated a 16 megabyte virtual address space.

3) Mapped files

Files are not accessed via a procedural interface (read record, write record, etc.), but by being associated with a range of virtual addresses and accessed as if they were memory.

4) Controlled sharing of information

The system organises sharing of program and data without imposing any restrictions or red tape on the user.

5) Transparent memory hierarchy

The system operates a three level memory hierarchy but the user is only aware of files and virtual addresses.

6) Minimal user constraints

The system attempts to restrain the user as little as possible in his use of files, languages, virtual addresses etc. There is a set of facilities provided by a standard subsystem, but the user may ignore it and easily provide his own subsystem if he so wishes.

7) Minimal information loss through crashes

If the hardware crashes an attempt is made to minimize the information loss.

8) No system degradation under load

The system cannot allocate more resources than it possesses, but it should allocate as much as possible to the users at all times. No thrashing.

9) Repeatability and enforced fairness

The resources used by a job should be dependent on its requirements, and not on other demands on the system; therefore if a job is run again it should use the same resource. A user should get a fair share of the system determined by his own requirements and the number of users of the system.

DESIGN PHASE

The design phase was unlike the corresponding stage of any other project known to the authors. It consisted of 15 months of academic discussion between people who were sure that they could write a much better operating system than the one supplied. This stage was followed by a very short period of ruthless decision taking. Of the major decisions taken at this time only one was seriously wrong and that one proved relatively easy to reverse. In spite of this fairly high success rate we would hesitate to suggest that others follow this procedure. The major constraint on the design can be seen at a glance from the diagram of EMAS structure (Figure 4). The Kernel and Director comprise about 144Kb of code out of a total of nearly 1Mb of system code plus all the user written code and packages. In view of the small number of people available and the limited timescale the only feasible strategy was to implement the Kernel and Director so that most of the rest of the code could be transferred by recompilation with minor amendments. Any such reimplementations must take into account the substantial hardware differences detailed in Table 1.




Figure 4: EMAS 4-75 Component Sizes

| 4-75 | 2900 |
|---|---|
| Register Orientated | Stack Orientated |
| Hand Coder's machines | Compiler Writer's machine |
| Large Pages (4 kb) | Small Pages (1kb) |
| Small VM (16Mb) | Large VM (8192Mb) |
| Small Store (1Mb max) | Vast Stores |
| Partial Paging (Non-priv only) | Fully Paged |
| Devices on Channels | Peripheral Controllers |
| Large Drum/Disc speed differential | Smaller Drum/Disc speed differential |
| Poor Comms support | Poor Comms Support |

Table 1: Comparison of 4-75 and 2900 Hardware

Against this background the original system was reviewed and revised to make the minimum changes necessary; the results of this review were as follows:-

a) System Structure

The basic structure of EMAS consisted of a small resident Supervisor and a paged supervisor process (Director) which together provided for each user a large virtual memory or address space. The low address part of the virtual memory is not accessible to the user and the various parts of the software are linked together by a message passing system and dispatcher¹³. These features have been transferred to EMAS 2900 with only one minor change: Director has been implemented as a set of privileged procedures running in the user's virtual machine, as in MULTICS, rather than a separate process associated on a one for one basis with each user process. It was necessary to use a separate process on the 4-75 as the protection mechanism could not otherwise prevent the user from corrupting Director. The excellent access control and System Call features of 2900 architecture have enabled the simpler and more efficient approach to be adopted.

b) File System

Central to EMAS is its File Store⁸ together with its Backup and Archiving System¹⁴. This holds named files for all users, each file consisting of an unstructured sequence of bytes of arbitrary length. The File System is virtual in that files are accessed by becoming part of the virtual memory (connection) rather than by any record access mechanism. Director maintains the file store and has ingenious algorithms for allocation and deallocation of space to avoid fragmentation. EMAS 2900 has adopted the EMAS file system in its entirety, even maintaining the unit of disk allocation at one EMAS

segment (64Kb) rather than one 2900 segment (256Kb). Having the unit of allocation at less than a segment involved small changes to the virtual memory management but this was preferred to a much larger unit of allocation which would cause fragmentation losses that would be unacceptable on the smaller discs (<100Mb). Further a consistent block size simplified movement of the file system maintenance utilities and facilitated a common archive (tape) store for both machines.

c) Configuration Independence

In EMAS, which was a prototype system, the configuration was built into the code. Flexibility was restricted to enabling the Operator to mark unserviceable devices as not available. In EMAS 2900 all configuration information appears in a resident read only segment which can be considered part of the machine on which the system runs. This segment could be nominated at load time by the Operator (a procedure used by ICL on 2950 and 2956 machines) but the design aim was to construct this segment by means of a General Reconnaissance Of Peripheral Equipment (GROPE). This would involve a two part system load. Initially a cut down supervisor is loaded. This performs the GROPE, constructs the read only segment, establishes communication with the Operator and loads the system. The design aim was to allow the system to run on any hardware configuration then announced without change.

d) Local Control

An important principle of EMAS was that of local decision-making. Decisions - particularly those affecting page replacement - are taken on information

relating to the behaviour of the relevant process and not on the basis of all processes currently in the multiprogramming set. This ensures reproducibility of program behaviour and prevents a program that exhibits anti-social paging behaviour from affecting other well-behaved processes. The rationale and effects of local control are discussed at length elsewhere¹⁰. However this organisation is not readily apparent in the original system which consisted of modules partitioned by function. These modules can make local decisions by reference to the process master segment¹³ (a sort of run-time database) of fixed and complex structure.

The design aim for the new system was to provide a module - the Local Controller - on a one for one basis with each process, to run in the process's address space. This method would reflect more accurately the structure of the system and would mean that the run-time database could consist of ordinary IMP variables and arrays. Thus the database layout would be the concern of the compiler, not the programmer, and arrays could be dynamic based on information discovered at GROPE time. The database would still page with the process and would not be in main store when the process was in the wait state. These advantages, particularly that of avoiding a fixed and unalterable database format, were considered to outweigh a quite serious disadvantage. 2900 architecture⁴ provides only one software stack switching interrupt, the OUT interrupt, which would have to be used for Director - Local Controller communication and also for Local Controller - Supervisor communication. Such double use is possible, but expensive, since the Interrupt Steering Table is constantly updated.

e) Virtual Memory Control

Great thought was given to the design of virtual memory management software

and tables since system overhead was mostly directly related to the amount of paging traffic. In those considerations much weight was given to various performance measurements made on the original EMAS system^{1,2,9}.

The first major decision concerned page size. System 4 had a 4K page while 2900 had 1K pages which could be grouped by software into extended pages (epages). The performance measurers were certain that a 1Mb 4-75 would perform best with 2K pages but software changes would be minimised by keeping to a 4K epage. The decision was to program with a potentially variable epage size but to start with a 4K page.

Secondly the EMAS memory hierarchy was critically scrutinised. In this each process has a core working set which is a true subset of the larger Active store (Drum)* working set which is itself a subset of the process's virtual memory. If no drums are serviceable there is no interactive service; this was reasonable on System 4 where the Discfiles are very slow but unreasonable on 2900. EMAS 2900 must be able to run without any fixed head paging device and also to run effectively with much less than the optimal amount of such storage.

Thirdly, the paging strategy was examined. Should EMAS 2900 use the EMAS method of preloading a process's working set of pages when that process enters the multiprogramming set (often called the Swopped Working Sets or SWS strategy) or should the simpler demand paging strategy be adopted? SWS was markedly superior on EMAS and has recently acquired powerful theoretical support⁵ so the decision here was to retain the SWS strategy.

* The paging device is referred to throughout as a Drum, although on 2900 it is in fact a small fixed-head disc.

Fourthly, program and data sharing was reconsidered in the light of 2900 hardware assistance - viz. a second (Public) segment table for all processes and also cascaded indirection of segment table entries. Although System 4 provided no hardware assistance, EMAS provided extensive program and data sharing in a manner transparent to users. The sharing was a most successful feature in that it enabled the operating system to effect substantial economies in core and drum space as well as a more modest saving in page transfers. These savings are particularly valuable since they increase with load; thus sharing helps performance most when the system is overloaded. Sharing via public segments is obviously possible but such sharing introduces a problem of selection; using indirect local segment table entries to point at a shared segment table does not remove the selection problem. Both methods have the serious side effect of combining the page use markers for shared material, thus preventing the accurate maintenance of the process's working set. EMAS 2900 ignores these hardware aids and uses the public segment table only as a local segment table for the Supervisor.

Lastly a minor change in the handling of unused page frames was incorporated. A quirk of implementation had resulted in EMAS losing the contents of a page frame when it had been successfully written out after use. If the page was wanted again it could not be extracted from the list of unused pageframes on those occasions when the frame had not yet been re-used. This inconsistency was never removed since the free page list had typically 20-50 entries as against a total of 1500 - 2000 "active" pages and it seemed likely that the gains from "recapturing" the page would be small. However in view of the large store sizes available on 2900 series, EMAS 2900 would recapture pages.

f) Communications

The communications on EMAS had been overtaken by events during the eight year life of the system. The design used ICL's hardware multiplexor (The Multi-Channel Communications Control Unit) which necessitated fixed resident buffers and a "Buffer Manager". Later the M.C.C.C.U. was replaced by a Front End Processor but the essentials of the software design were unchanged. The design worked well when all terminals were operating at ten characters per second. What had not been foreseen was the great increase in terminal speed and hence I/O in a terminal session; a further surprise was the attraction of EMAS's secure file system. Programmers using machines in Newcastle, Cambridge, Harwell and Manchester kept their programs and data on EMAS and transmitted them daily. The amount of paging traffic on EMAS concerned with trivial buffer filling and emptying rapidly became insufferable. A radical rethink was required, complicated by the inability of ICL to provide a communications device that seemed likely to fulfil the demands EMAS 2900 would make on it.

The elegant solution adopted will be fully described in a subsequent paper. The main features are as follows:

- 1) A Digital PDP11 as Front End Processor interfaced to ICL's 2900 Application Module interface via locally built hardware.
- 2) All control functions to reside in FEP.
- 3) Transfers between processes and the FEP to be made using large buffers in virtual memory.

Two features would enable the virtual buffering to be incorporated. Firstly the address translation facility of Peripheral Controllers; secondly the software structure of EMAS 2900 necessary for local control. Just as a Local Controller manages the paging of the working set of a user process, so a modified Local Controller - the Communications Controller - can manage the paging of all virtual buffers involved in transfers. Indeed, since both Local and Communication Controllers could use the full range of Supervisor services, it seemed likely that very little new code would be required. The data sharing of EMAS 2900 would operate on such buffer pages without any amendment.

The designers were aware that they were adopting a novel and high risk communications strategy - alternatives, however, were uninviting.

TOOLS

A brief description of the tools used in developing EMAS 2900 seems necessary. No new tools were constructed as part of the project - the most important were already available and the others were transferred at an early stage from the original system to the new.

The least sung but most valuable tool was the original EMAS system in spite of being severely overloaded. Its terminal access, editors, compilers and debugging features were invaluable. Above all its secure file store preserved and protected the source code. No source code was lost or destroyed from the start of the project in October 1976 until January 1978 when development was moved to the new system. In spite of the load on the system it was usually possible to edit a module, recompile it, link a new

test supervisor and write it to magnetic tape for testing on the 2970 in about half an hour of elapsed time.

The second tool was the same dialect of the Edinburgh Implementation Language (IMP)¹² as was used in the original system. We do not claim that IMP is technically superior to more modern languages such as CONCURRENT PASCAL⁶, but there were compelling reasons to use it. Firstly it existed - the language and three compilers (for System 4, 2900 and also a System 4 to 2900 cross compiler), together with a large selection of source code formatters, diagnostic packages, etc. Secondly, all the implementers understood every nuance of the language and thirdly, we wished to transport source code from the old system.

In retrospect two aspects of IMP were particularly valuable. In the early stages the ability to obtain source language diagnostics from any failure saved many hours of debugging time. Later the ability to find out how often each source statement had been executed enabled the small amount of effort available for tuning to be deployed effectively.

The third tool available was system monitoring code. Much effort had been spent on monitoring the early EMAS system but two particular areas had proved of lasting value. The ability to monitor the message passing mechanism in a highly selective way has proved invaluable in disentangling curious interactions between components. Also of value was the ability to write the system to magnetic tape and later examine the tape with a utility which could reconstruct, in source terms, complex tables and linked lists.

Fourthly there was the Edinburgh Remote Terminal Emulator - ERTE¹. This could produce between 8 and 64 interactive terminals for testing, performance measurement and benchmarking.

IMPLEMENTATION

The EMAS 2900 project was informal in structure, short in duration and small in size. It began in October 1976 with the aim of providing a service to invited users in January 1978 and a full service in October 1978. The project started with the four authors augmented by a group of six or seven individuals who made specific contributions in an area of expertise (e.g. device handlers). The contributions of this second group were limited in time. The project increased in size when the system was able to support superstructure and had reached eight full-time and four or five part-time members by September 1978. The total effort applied up to the start of full service was about 15 man years - very close to the prediction of Whitfield and Wight¹³. During the two year project about 1000 hours of dedicated 2970 time was used. This total does not include terminal access to the initial user service or any time used on System 4.

The implementation was mostly free from unpleasant surprises. GROPEing was very difficult to implement since it involved using the hardware at a basic level, where there were great differences between the various machines. Had we foreseen the pain involved or the changes needed when the hardware was upgraded - a fairly frequent occurrence in the early years of the 2900 series - we would not have attempted this extension. Nevertheless now that all is working and understood, it greatly simplifies software maintenance to have a single self-configuring software product.

The early communication software was shaken by the discovery of a curious and then undocumented omission in all 2900 Peripheral Controllers. These Controllers perform address translation with full protection just as does the processor - this is a very desirable hardware feature. Having done all

this the controllers fail to update the read and write history bits in the page tables so that the Local Controller could miss seeing that the virtual communications buffer had been used or updated. A software bodge managed to cure this but at the cost of writing out pages that have not been updated.

In spite of these minor tribulations the implementation timescales were met. Although the initial user service was not started until May 1978 the full service started in October as planned with rather more than the promised facilities. The final code size of the Supervisor was greater by about 30% than for the original system. Almost all the increase was in driving the more sophisticated 2900 peripherals. The drum driver was markedly more complicated since the 2900 Sector File Controller operates with 1K sectors. Each EMAS 2900 transfer involves initiating four requests and fielding four terminations. The data areas have also increased in size partly because of the larger VM and segment sizes on 2900 but mostly because the absence of 16-bit integers has resulted in 32-bit entries being required. Nevertheless, about 200 4K page frames are available for user pages out of 1Mb store on both systems. This is partly because there are no communications buffers in EMAS 2900 and also because the local segment table is now part of the Local Controller stack and pages with the process.

PERFORMANCE

The user view of EMAS 2900 performance is very satisfactory. With 32 terminals, the 2970 provides a much livelier service than the 4-75. It is necessary to have about 45 terminals active before the service degrades to a level that will not pass the interactive part of the Glasgow benchmark.

However, since the 2970 is about twice as fast as the 4-75 it should be expected to support 64 terminals. It is instructive to compare the EMAS 2970 timing measurements with the figures for 4-75 published by Whitfield and Wight¹³. Exact comparison is difficult but it is clear that a higher proportion of time is used in the Supervisor on 2900 than on 4-75. It is also clear that device driving is responsible for much of this. In Whitfield and Wight, 1 million drum transfers have taken 973 secs - i.e. about 970 micro secs of 4-75 CPU time each, whereas in Table 2 314000 drum transfers and corresponding interrupts have used 514 secs, i.e. about 1630 micro secs of 2970 time or 3260 micro secs of 4-75 time per transfer.

The other revelation from Table 2 is that more than 50% of page exception interrupts can be satisfied without a transfer - 16% because the page is being used by another process and 37% by recapturing the page from the free list. These figures show the advantages of abandoning paging devices and using very large main memories.

The virtual buffering technique used in communicating with interactive terminals is very successful and terminals of all speeds have been attached. A buffer of one page (4Kb) is sufficient for slow and medium speed terminals although the software will accept any size up to 64Kb. The virtual buffering is so trouble-free that the local printers and card readers have been attached as pseudo-remote devices. Using the maximum virtual buffer size it is possible for the spooling manager - a paged process - to read or print 64Kb without any intermediate transactions. This is a great improvement on the original system, which required the spool manager to page in every 4000 bytes read or printed. The other design aim of the communications software - that of protocol independence - seems to have been achieved judging by the experience of the University of Kent, who run EMAS 2900 on their ICL 2960. Kent operate a very different local

network from Edinburgh but they are able to run the EMAS 2900 software unchanged. All necessary alterations are confined to the FEP and these were accomplished and tested in a few months. The CPU overheads of the communications software at about 5% of the CPU on 2970 are higher than on 4-75, where they were held to 1.5% of the CPU. This price is well worth paying for the machine independence and reduction in paging.

Two desirable attributes have survived the re-implementation intact. The first is robustness - the mean time between software crashes on EMAS has always been high and it tended to infinity once development ceased. EMAS 2900, although still at an early stage, already survives hundreds of hours between software crashes. The second is the user interface as provided by the Standard Subsystem⁷. This feature together with compatible compilers makes transferring work from the old to the new system relatively painless.

We still regard EMAS 2900 as at an early stage in its life and expect the programme of tuning, measurement and analysis to yield a steady improvement in performance over the next two or three years. As part of these experiments it is hoped to try various page sizes in an experimental system.

BENCHMARKING

The presence in Edinburgh of the Remote Terminal Emulator (ERTE), used to establish the interactive part of Glasgow benchmark, proved very useful. The 32 terminal scripts were subdivided into groups of 8 and duplicated, thus enabling 8, 16, 24, 32, 40, 48, 56 or 64 terminals to be emulated. (Strictly it is not possible to subdivide the benchmark as all scripts are

different). These emulated users could be used for testing the infant EMAS 2900 system to find bugs and measure performance. Having eliminated a number of bugs ERTE was then used to time the system and immediately produced a very surprising result. EMAS 2900 performed as well or better with a simple demand paging algorithm as it did with the SWS algorithm which had been so successful on System 4 EMAS. This conclusion seemed true across a range of configurations, including those with no drums. The designers were extremely surprised but it proved easy to correct the faulty design decision and simplify the software by deleting the preloading of working sets. Tables 4 and 5 give figures measured by ERTE for benchmark runs. The technique used was to start the benchmark, wait ten minutes so that all terminals were active and then take measurements for a 30 minute window. This was much less time-consuming than running the entire benchmark. From such complete runs as were made it seemed necessary to achieve 50-55 commands/minute on the 32 terminal version and 100-110 commands/minute on the 64 terminal version if the benchmark was to be completed in the 2-hour period specified, the average amount of work done per simulated user being kept constant. When examining these tables one should realise:

- a) That the reaction time is measured from the last character of the command being typed until the first character of the reply.
- b) That the response time is measured from the last character of the command being typed to the first character of the next prompt reaching the terminal. Thus listing 1000 characters to a 10 character/sec terminal must have a response time of at least 100 seconds.
- c) That some commands in this benchmark compile and run very large programs of many thousand statements, and that a small number of such commands have a large effect on average response times.

Experiment Y shows that EMAS 2900 can support 64 terminals on a 2980 along with a lot of simultaneous batch work while maintaining a very good response - more than 90% of editing reactions were less than 2 secs. This is sufficient to pass the 64 terminal benchmark by a large margin.

Experiment B shows that a 2970 can pass the 32 terminal version of the benchmark (as specified at Southampton University) with good response when using 12Mb of drum. Experiment D shows that it is just possible to pass the Southampton benchmark with no drums at all but that response is sluggish with only 23% of editing reactions in 2 seconds.

EMAS 2900 SUP26A TIMING MEASUREMENTS

| Service | Calls | Time (secs) | Average (msecs) | % of total | Av.inst. /call | Notes |
|-----------------|---------|----------------|--------------------|---------------|-------------------|-------------------------|
| IDLE TIME | 82723 | 1090.121 | 13.178 | 9.2% | 1 | Idle time* |
| NOWORK TIME | 64528 | 4662.809 | 72.260 | 39.5% | 1 | |
| DEADLOCK RCVRY | 140 | 0.243 | 1.733 | 0.0% | 1154 | Paging |
| SCHEDULE | 65969 | 90.130 | 1.366 | 0.8% | 609 | |
| PAGETURN | 1082650 | 703.670 | 0.650 | 6.0% | 241 | |
| GET EPAGE | 1574 | 1.334 | 0.848 | 0.0% | 322 | |
| RETURNEPAGE | 385587 | 60.661 | 0.157 | 0.5% | 72 | |
| SEMAPHORE | 16 | 0.014 | 0.866 | 0.0% | 323 | Drum loading |
| ACTIVEMEM | 54980 | 51.455 | 0.936 | 0.4% | 380 | |
| ELAPSEDINT | 11623 | 14.256 | 1.277 | 0.1% | 686 | |
| UPDATE TIME | 5636 | 6.733 | 1.195 | 0.1% | 566 | |
| DPONPUTONQ | 1 | 0.001 | 0.504 | 0.0% | 214 | Disc driving |
| TURNONEP | 1127 | 0.332 | 0.295 | 0.0% | 141 | |
| ACTIVEMEM(POLL) | 1127 | 1.570 | 1.393 | 0.0% | 672 | |
| SCHEDULE(OPER) | 2254 | 13.252 | 5.880 | 0.1% | 2854 | |
| DISC | 3766 | 2.511 | 0.667 | 0.0% | 331 | |
| DISC TRANSFERS | 215211 | 113.373 | 0.527 | 1.0% | 227 | File system maintenance |
| DISC INTERRUPT | 138060 | 194.097 | 1.406 | 1.6% | 584 | |
| MOVE/1 | 1030 | 2.070 | 2.009 | 0.05 | 1065 | Drum driving |
| MOVE/2 | 6507 | 3.772 | 0.580 | 0.0% | 262 | |
| DRUM TRANSFERS | 314520 | 289.527 | 0.921 | 2.5% | 369 | Slow peripherals |
| DRUM INTERRUPT | 172526 | 224.787 | 1.303 | 1.9% | 499 | |
| GPC | 100433 | 84.977 | 0.846 | 0.7% | 369 | Communications |
| TAPE | 2869 | 4.795 | 1.671 | 0.0% | 765 | |
| OPER | 13663 | 12.341 | 0.903 | 0.1% | 420 | |
| LP ADAPTOR | 1567 | 15.445 | 9.856 | 0.1% | 716 | |
| CR ADAPTOR | 4 | 0.008 | 1.998 | 0.0% | 1274 | |
| PRINTER | 5811 | 2.721 | 0.468 | 0.0% | 242 | Working set maintenance |
| GPC INTERRUPT | 126689 | 143.178 | 1.130 | 1.2% | 422 | |
| COMMS CONTROL/1 | 29341 | 22.507 | 0.767 | 0.2% | 282 | Working set maintenance |
| COMMS CONTROL/2 | 91621 | 68.477 | 0.747 | 0.6% | 310 | |
| FEP ADAPTOR | 180115 | 164.533 | 0.913 | 1.4% | 376 | Working set maintenance |
| LOCAL CONTROL | 212570 | 473.558 | 2.228 | 4.0% | 923 | |
| FOREGRND USERS | 598005 | 2929.841 | 4.899 | 24.8% | 0 | |
| BACKGRND USERS | 17034 | 161.125 | 9.459 | 1.4% | 0 | |

INTERRUPT/ACTIVATE ETC. = 206.897 SECS (1.8%)

DRUMSIZE = 3072

OVERALLOC = 40

RECAPTURES = 37%

SHARED PAGES = 16%

SOFTWARE INWARD CALLS = 0

* NOWORK TIME is idle time when all users are voluntarily in the wait state.
IDLE TIME is all other idle time.

Table 2: 2970 Performance Measurements, 4/1/1980
(Optimised Supervisor)

2900 SUP26A TIMING MEASUREMENTS

| Service | Calls | Time (secs) | Average (msecs) | % of total | Av.inst. /call | Notes |
|-----------------|---------|----------------|--------------------|---------------|-------------------|----------------------------|
| IDLE TIME | 94951 | 927.211 | 9.765 | 16.3% | 1 | Idle time* |
| NOWORK TIME | 30351 | 637.087 | 20.991 | 11.2% | 1 | |
| DEADLOCK RCVRY | 0 | | | | | Paging |
| SCHEDULE | 103198 | 43.426 | 0.421 | 0.8% | 866 | |
| PAGETURN | 1669349 | 305.307 | 0.183 | 5.4% | 333 | |
| GET EPAGE | 497 | 0.117 | 0.235 | 0.0% | 418 | |
| RETURNPAGE | 520624 | 29.811 | 0.057 | 0.5% | 115 | |
| SEMAPHORE | 36 | 0.008 | 0.221 | 0.0% | 351 | Drum loading |
| ACTIVEMEM | 92184 | 24.882 | 0.270 | 0.4% | 509 | |
| ELAPSEDINT | 5898 | 2.514 | 0.426 | 0.0% | 737 | |
| UPDATE TIME | 2709 | 1.093 | 0.403 | 0.0% | 815 | |
| DPONPUTONQ | 2 | 0.000 | 0.140 | 0.0% | 248 | |
| TURNONEP | 541 | 0.057 | 0.105 | 0.0% | 175 | Disc driving |
| ACTIVEMEM(POLL) | 541 | 0.302 | 0.559 | 0.0% | 1181 | |
| SCHEDULE(OPER) | 1083 | 2,357 | 2.177 | 0.0% | 4629 | |
| DISC | 1814 | 0.422 | 0.232 | 0.0% | 434 | |
| DISC TRANSFERS | 329407 | 43.740 | 0.133 | 0.8% | 228 | |
| DISC INTERRUPT | 193225 | 69.728 | 0.361 | 1.2% | 663 | File system maintenance |
| MOVE/1 | 2016 | 1.915 | 0.950 | 0.0% | 1922 | |
| MOVE/2 | 18351 | 3.079 | 0.168 | 0.1% | 306 | Drum driving |
| DRUM TRANSFERS | 368374 | 82.430 | 0.224 | 1.5% | 340 | |
| DRUM INTERRUPT | 181686 | 76.109 | 0.419 | 1.3% | 648 | Slow peripherals |
| GPC | 126537 | 33.920 | 0.268 | 0.6% | 536 | |
| TAPE | 12769 | 130.639 | 10.231 | 2.3% | 32454 | |
| OPER | 13239 | 12.191 | 0.921 | 0.2% | 2426 | |
| LP ADAPTOR | 4646 | 11.291 | 2.430 | 0.2% | 847 | |
| CR ADAPTOR | 27 | 0.080 | 2.979 | 0.0% | 8467 | Communicat- ions |
| PRINTER | 5058 | 1.430 | 0.283 | 0.0% | 526 | |
| GPC INTERRUPT | 172615 | 59.786 | 0.346 | 1.1% | 611 | |
| COMMS CONTROL/1 | 37839 | 8.637 | 0.228 | 0.2% | 382 | Working set maintenance |
| COMMS CONTROL/2 | 133658 | 32.780 | 0.245 | 0.6% | 429 | |
| FEP ADAPTOR | 242966 | 69.579 | 0.286 | 1.2% | 523 | |
| LOCAL CONTROL | 250885 | 253.864 | 1.012 | 4.5% | 1659 | |
| FOREGRND USERS | 805612 | 1374.458 | 1.706 | 14.2% | 0 | |
| BACKGRND USERS | 231967 | 1396.497 | 6.020 | 24.6% | 0 | |

INTERRUPT/ACTIVATE ETC. = 43.168 SECS (0.8%)

DRUMSIZE = 4608

OVERALLOC = 113

RECAPTURES = 39%

SHARED PAGES = 28%

SOFTWARE INWARD CALLS = 0

* NOWORK TIME is idle time when all users are voluntarily in the wait state.
IDLE TIME is all other idle time.

Table 3: 2980 Performance Measurements, 15/1/1980
(Diagnostic Supervisor)

2970 1Mb Core 4 EDS100s on 1DFC

| Experiment | A | B | C | D |
|----------------------|------|------|------|------|
| Mbytes of Drum | 12 | 12 | 12 | 0 |
| Batch Streams | 0 | 0 | 0 | 0 |
| Interactive users | 16 | 32 | 47 | 32 |
| All commands | | | | |
| average reaction | 2.2 | 7.2 | 13.9 | 15.8 |
| average response | 10.5 | 15.4 | 22.9 | 23.6 |
| commands/min | 40 | 64 | 71 | 53 |
| Editing | | | | |
| average reaction | 0.5 | 1.4 | 2.5 | 4.9 |
| average response | 3.7 | 4.9 | 6.5 | 8.2 |
| % reactions < 2 secs | 84 | 65 | 47 | 23 |
| CPU usage | | | | |
| % in supervisor | 22.5 | 42.0 | 43.2 | 36.2 |
| % in user | 37.7 | 56.5 | 51.7 | 46.4 |
| % idle + no work | 39.8 | 1.5 | 5.1 | 17.4 |
| Paging | | | | |
| % Recaptured | 43 | 27 | 23 | 16 |
| % Shared | 19 | 20 | 20 | 34 |

Table 4: EMAS 2970 Benchmarking (31.1.79 and 3.3.79)

| Experiment | X | Y |
|----------------------|------|------|
| Batch Streams | 8 | 8 |
| Interactive users | 48 | 63 |
| All commands | | |
| average reaction | 3.0 | 4.3 |
| average response | 10.2 | 11.6 |
| commands/min | 118 | 145 |
| Editing | | |
| average reaction | 0.7 | 0.9 |
| average response | 4.2 | 4.5 |
| % reactions < 2 secs | 97 | 92 |
| CPU usage | | |
| % in supervisor | 22.0 | 25.4 |
| % in user | 78.0 | 74.6 |
| % idle + no work | 0.0 | 0.0 |
| Paging | | |
| % Recaptured | 29 | 26 |
| % Shared | 43 | 45 |

Table 5: EMAS 2980 Benchmarking (6.4.79)

CONCLUSION

The relatively uncommon practice of re-implementing a system has produced a major operating system using a fraction of the effort that was required to produce the original software or other comparable systems. The self-restraint of the designers in restricting changes to the minimum has resulted in the new system retaining the user interface, reliability and high performance of the original. With the use of relatively machine-independent high-level languages re-implementation may become more widely considered as a technique for moving to new hardware with minimum disruption.

AVAILABILITY

EMAS 2900 is currently running on 2960, 2970 and 2980 machines. It has run on a 2976 and will soon be mounted on a dual processor 2972. It will not run on a 2950 or 2956. The University will make the software available in source code form on request. There will be no charge for this in the case of Educational or Research institutions. The University is also able to supply the special hardware developed for the communications interface.

ACKNOWLEDGEMENTS

We would like to thank Professor H. Whitfield and the other designers of EMAS for producing a system that it has been a pleasure to re-implement. The following, among others, contributed ideas, suggestions and that most precious commodity, reliable, tested code: A. Anderson, W. Laing, C.D. McArthur, R.R. McLeod, J. Maddock, P. Robertson, A. Shaw, F. Stacey and A. Waller. B.A.C. Gilmore designed and implemented the FEP software. We also thank the Post Office Data Processing Service for seconding Messrs. J. Maddock and A. Waller to work with us.

REFERENCES

1. J.C. Adams, W.S. Currie and B.A.C. Gilmore, 'The Structure and Uses of the Edinburgh Remote Emulator', Software Practice and Experience, 8, 451-459 (1978).
2. J.C. Adams and G.E. Millard, 'Performance Measurement on the Edinburgh Multi-Access System', Proc. Int. Computing Symposium, Antibes, 105-112 (1975).
3. J.C. Adams, 'Performance measurement and evaluation of time shared virtual memory systems', Ph.D Thesis, Edinburgh University (1977).
4. J.K Buckle, 'The ICL 2900 Series', Macmillan, 1978.
5. P.J. Denning, 'Working sets past and present', Purdue University Report CSD-TR-276 (1978).
6. P.B. Hansen, 'The programming language Concurrent Pascal', IEEE Transactions on Software Engineering, 1, 2 199-207 (1975).
7. G.E. Millard, D.J. Rees and H. Whitfield, 'The Standard EMAS Subsystem', Computer J. 18, 213-219 (1975).
8. D.J. Rees, 'The EMAS Director', Computer J. 18, 122-130 (1975).
9. A.L. Schoute, 'Modelling of virtual memory computer systems with multiple classes of process', Ph.D Thesis, University of Groningen (1978).

10. N.H. Shelness, P.D. Stephens, and H. Whitfield, 'The Edinburgh Multi-Access System, Scheduling and Allocation Procedures in the Resident Supervisor', RAIRO (Informatic, Computer Science) B, 3, 29-45 (1975).
11. J.G. Sime, 'Benchmark for the ICL 1906S Computer System', University of Glasgow, 1973.
12. P.D. Stephens, 'The IMP Language and Compiler', Computer J., 17, 216-223 (1974).
13. H. Whitfield and A.S. Wight, 'EMAS - The Edinburgh Multi-Access System', Computer J., 18, 331-346 (1973).
14. A.S. Wight., 'The EMAS Archiving Program', Computer J., 18, 131-134 (1975).