

**South African  
Computer  
Journal  
Number 21  
August 1998**

**Suid-Afrikaanse  
Rekenaar-  
tydskrif  
Nommer 21  
Augustus 1998**

**Computer Science  
and  
Information Systems**

**Rekenaarwetenskap  
en  
Inligtingstelsels**

**The South African  
Computer Journal**

*An official publication of the Computer Society  
of South Africa and the South African Institute of  
Computer Scientists*

**Die Suid-Afrikaanse  
Rekenaartydskrif**

*'n Amptelike publikasie van die Rekenaarvereniging  
van Suid-Afrika en die Suid-Afrikaanse Instituut  
vir Rekenaarwetenskaplikes.*

**World-Wide Web: <http://www.cs.up.ac.za/sacj/>**

**Editor**

Prof. Derrick G. Kourie  
Department of Computer Science  
University of Pretoria  
Hatfield 0083  
dkourie@cs.up.ac.za

**Sub-editor: Information Systems**

Prof. Niek du Plooy  
Department of Informatics  
University of Pretoria  
Hatfield 0083  
nduplooy@econ.up.ac.za

**Production Editor**

John Botha  
Department of Computer Science  
University of Pretoria  
Hatfield 0083  
sacj\_production@cs.up.ac.za

**Editorial Board**

Prof. Judith M. Bishop  
University of Pretoria, South Africa  
jbishop@cs.up.ac.za

Prof. R. Nigel Horspool  
University of Victoria, Canada  
nigelh@csr.csc.uvic.ca

Prof. Richard J. Boland  
Case Western University, U.S.A.  
boland@spider.cwrv.edu

Prof. Fred H. Lochovsky  
University of Science and Technology, Hong Kong  
fred@cs.ust.hk

Prof. Ian Cloete  
University of Stellenbosch, South Africa  
ian@cs.sun.ac.za

Prof. Kalle Lyytinen  
University of Jyväskylä, Finland  
kalle@cs.jyu.fi

Prof. Trevor D. Crossman  
University of Natal, South Africa  
crossman@bis.und.ac.za

Dr. Jonathan Miller  
University of Cape Town, South Africa  
jmiller@gsb2.uct.ac.za

Prof. Donald. D. Cowan  
University of Waterloo, Canada  
dcowan@csg.uwaterloo.ca

Prof. Mary L. Soffa  
University of Pittsburgh, U.S.A.  
soffa@cs.pitt.edu

Prof. Jürg Gutknecht  
ETH, Zürich, Switzerland  
gutknecht@inf.eth.ch

Prof. Basie H. von Solms  
Rand Afrikaanse Universiteit, South Africa  
basie@rkw.rau.ac.za

**Subscriptions**

	Annual	Single copy
Southern Africa	R50.00	R25.00
Elsewhere	US\$30.00	US\$15.00

An additional US\$15 per year is charged for airmail outside Southern Africa

to be sent to:

*Computer Society of South Africa  
Box 1714, Halfway House, 1685  
Phone: +27 (11) 315-1319 Fax: +27 (11) 315-2276*

# Cultivating the information systems discipline

Niek du Plooy, Sub-Editor: Information Systems

Whether by 'information system' we mean a simple bookkeeping system for a small business, or a monolithic integrated 'management information system' for a global corporation, all organisations currently need information systems in order to function effectively. The computer and business community at large have readily adopted and accepted the use of the term 'information systems', but perhaps without too much real thought being given to any more profound meaning of the term. Departments bearing that name (or something very similar) are commonly found in organisations. But can the same be said for the 'academic' use of the term, as in describing the information systems discipline? Has it been 'accepted' as a separate scientific discipline?

The term 'discipline' is often loosely applied to indicate the scientific 'field', that is, the organised 'body of knowledge' or 'domains of discourse' within which (mainly) academic activities concerning a specific topic or a number of related topics, are conducted. [3] point out that a scientific discipline has a certain paradigm associated with it, meaning that researchers in that discipline are familiar with the research topics, the research methods and the accepted ways to interpret the results in their chosen field. A discipline is further strengthened and consolidated by the educational process whereby a researcher becomes a practitioner in that discipline, initially through the pursuit of academic degrees and thereafter, through recognition amongst his/her peers. Formal study in a particular discipline results in the value sets and exemplars (the 'paradigm') of that discipline being adopted by the student, either consciously or unconsciously.

Is 'information systems' truly a recognised scientific discipline such as this? In the past, prominent authors such as Peter Keen did not think so [15, 16]. He deplored the lack of a cumulative tradition and advocated that one be built up, asking for a clarification of the reference disciplines of this new science and a definition of its dependent variable and the building of a cumulative tradition, amongst other things. [1] however, disputed Keen's position and pointed to strong links between research and practice found in their analysis. [11] showed clearly that 'orthodoxy' exists in many aspects of information systems, i.e. in information systems methodologies as well as in other areas of information systems development. This claim was supported by [13] who, in a detailed study based on papers in scientific journals, scientific conferences and textbooks, identified seven different but complementary 'schools of thought' within the field of information systems. In a study of leading universities and

leading researchers in decision support systems, [9] provide exemplars, at least for that particular sub-discipline. [5] conclude from a citation study of journal influence during the period 1981 through 1985, that the discipline of information systems has attained stability and that it is in no danger of dying. It seems therefore, that Keen's despair is unfounded and that information systems have indeed grown into a separate, identifiable discipline, even if the field is best described as a 'fragmented adhocracy' ([3]).

The existence of an established scientific community in information systems has been given formal recognition by the recent formation of the Association for Information Systems, a professional society in the tradition of scientific societies, with 1400 members in 35 countries. A recently compiled directory of information systems academics contains entries on some 4,500 researchers from more than 1,000 institutions. A number of basic University and other curricula for information systems education have been published over the years [2, 6, 18]. The most recent of these is Curriculum '95, a joint effort by the ACM, AIS, DPMA, IAIM and ICIS [10, 7]. The most popular discussion group on the Internet (ISWorldNet) devoted exclusively to information systems matters has a membership which in 1997 approached 1829 from 53 countries [14]. A well-defined scientific community therefore exists.

In addition, if the existence of sound academic scholarship is further testimony to the existence of a 'discipline', then information systems can proudly point towards a dramatic growth over the past three decades in the number of scientific journals reporting on research in this area [12]. An even more recent study on research outlets showed that, amongst twenty-seven established journals carrying articles in this field, at least three of the most highly rated top ten are devoted exclusively to the discipline.

Yet, can it be said that the information systems discipline has been conclusively defined and that the research problems and research methodologies prescribed for it have been accepted by all who consider themselves to be working in this field? A re-examination and extension of an earlier (1988) list of keywords for use in classifying information systems literature [4] includes a list of the reference disciplines of information systems, as well as lists of the external environment, the technology, the organisational environment, etc., of information systems. We could argue that this very comprehensive list of keywords (nearly 1300) and other classifications define and describe the discipline of information systems accurately and usefully. For instance, the reference disciplines were listed as:

behavioural science, computer science, decision theory, information theory, organisation theory, management theory, language theories, systems theory, research, social science, management science, artificial intelligence, economic theory, ergonomics, political science, psychology. This list reflects the interdisciplinary or pluralistic nature of information systems.

In the same vein, [19] did a study on the themes of submissions to the journal *Information Systems Research* and produced a list of keywords, concepts and associations that characterise the categories into which they grouped the research questions of articles submitted. This list demonstrates conclusively that the subareas of the discipline (organisational, behavioural and managerial issues) are well established and attract a large number of researchers on a long-term basis. Swanson & Ramiller conclude by observing that the discipline still exhibits the 'fragmented adhoc-racy' identified by Banville & Landry, and is still topically diverse and '...based on appeals to significantly different and partly incommensurate reference disciplines'.

Thus, fragmentation can have adverse effects – something that information systems researchers should be aware of. However, fragmentation of the discipline of information systems may be evident in the field for a very long time. As has been pointed out [17, 8], the discipline as a whole follows trends in information technology, and researchers tend to build their interests around new technology (e.g. the earlier interest in expert systems and decision support systems, and current interest in computer-supported co-operative work). As information technology evolves, so the research interests will follow these new directions. Although we may wish it were different, it remains a fact that information technology is still a major reference discipline of information systems, and will remain so as long as researchers struggle to separate the fundamental or common issues in different fields from the technological ones.

Clearly, then, information systems is internationally well-established as a flourishing discipline. In the Southern African context it is important that the discipline should not merely flourish but be seen to flourish. To this end, this editorial calls on academics and especially on practitioners to add your contributions, via a submission to SACJ.

## References

- [1] M. Alavi and P. Carlson. A review of mis research and disciplinary development. *Journal of Management Information Systems*, 8(4):45–62, 1992.
- [2] R.L. Ashenhurst. Curriculum recommendations for graduate professional programs in information systems. *Communications of the ACM*, 15(5):364–398, 1972.
- [3] C. Banville and M. Landry. Can the field of mis be disciplined? *Communications of the ACM*, 32(1):48–60, 1989.
- [4] H. Barki, S. Rivard, and J. Talbot. A keyword classification scheme for is research literature: An update. *MIS Quarterly*, pages 209–226, June 1993.
- [5] R.B. Cooper and M. Blair, D. and Pao. Communicating mis research: A citation study of journal influence. *Information Processing & Management*, 29(1):113–127, 1993.
- [6] J.L. Couger. Curriculum recommendations for undergraduate programs in information systems. *Communications of the ACM*, 16(12):727–749, 1973.
- [7] J.L. et al. Couger. Guidelines for undergraduate is curriculum. *MIS Quarterly*, 19(3):341–360, 1995.
- [8] P. Ein-Dor and E. Segev. A classification of information systems: Analysis and interpretation. *Information Systems Research*, 4(2):166–204, 1993.
- [9] S.B. Eom and S.M. Lee. Leading us universities and most influential contributors in decision support systems research (1971-1989). *Decision Support Systems*, 9:237–244, 1993.
- [10] J.T. et al. Gorgone. Information systems '95 curriculum model — a collaborative effort. *Data Base*, 25(4):5–8, 1994.
- [11] R.A. Hirschheim, J. Iivari, and H. Klein. An assumption analysis of the five emergent approaches to information systems development (isd): A contribution to the paradigmatic foundations of alternative approaches to isd. Working paper, Binghamton, 1995.
- [12] C.W. Holsapple, L.E. Johnson, H. Manakyan, and J. Tanner. Business computing research journals: A normalized citation analysis. *Journal of Management Information Systems*, 11(1):131–140, 1994.
- [13] J. Iivari. A paradigmatic analysis of contemporary schools of is development. *European Journal of Information Systems*, 1(4):249–272, 1991.
- [14] B. Ives. Posting to listserv.hea.ieon 21 July 1997.
- [15] P.G.W. Keen. Mis research: Reference disciplines and a cumulative tradition. In *Proceedings of the First International Conference on Information Systems*, Philadelphia, pages 9–18, 1980.
- [16] P.G.W. Keen. Relevance and rigor in information systems research: Improving quality, confidence, cohesion and impact. In H-E. Nissen, H.K. Klein, and R.A. Hirschheim, editors, *Information Systems Research: Contemporary Approaches & Emergent Traditions*. North-Holland, Amsterdam, 1991.

- [17] K. Lyytinen. Information systems and critical theory. In M. Alvesson and H. Willmott, editors, *Critical Management Studies*. Sage Publications, London, 1992.
- [18] J. Nunamaker, J.L. Couger, and G.B. Davis. Information systems curriculum recommendations for the 80s: Undergraduate and graduate programs. *Communications of the ACM*, 25(11):781–805, 1982.
- [19] E.B. Swanson and N.C. Ramiller. Information systems research thematics: Submissions to a new journal, 1987-1992. *Information Systems Research*, 4(4):299–330, 1993.

# Preliminary Investigation of the RAMpage Memory Hierarchy

Philip Machanick<sup>a</sup>Pierre Salverda<sup>b</sup>

Department of Computer Science, University of the Witwatersrand,  
<sup>a</sup>philip@cs.wits.ac.za, <sup>b</sup>psalverd@cs.wits.ac.za

## Abstract

*The RAMpage memory hierarchy addresses the growing concern about the memory wall – the possibility that the CPU-DRAM speed gap will ultimately limit the benefits of rapid improvement in CPU speed. Reducing references to DRAM is an increasingly desirable goal as CPU speed improves relative to DRAM. As the cost of a DRAM reference increases, it makes increasing sense to consider options like pinning crucial parts of the operating system in at least the lowest-level cache, and to consider possibilities like context switches on references to DRAM. All these factors combine to make it increasingly desirable to treat DRAM as a paging device, while moving the main memory a level up to the lowest level of SRAM. The RAMpage hierarchy relegates DRAM to the role of a first-level paging device. Results presented here are for a preliminary simulation of the RAMpage hierarchy, and show that, if current memory system and CPU trends continue, the RAMpage strategy will become increasingly viable. Even with current miss costs and without implementing all features favourable to the RAMpage hierarchy, simulations show that it is possible to achieve run times up to 25% faster than those for a conventional hierarchy. Furthermore, RAMpage scales up better than the conventional hierarchy (as simulated) in that performance degrades less as DRAM reference costs increase.*

**Keywords:** memory hierarchy, caches, paging, software-controlled replacement

**Computing Review Categories:** B.3, C.4, D.4.2

## 1 Introduction

There has been much discussion in recent years of the “memory wall” [28, 17, 27] – a consequence of a growing CPU-DRAM speed gap [12, 5].

One approach to dealing with this growing CPU-DRAM speed gap is to focus on reduction of misses, even strategies which make each miss cost more. For example, past work on software-controlled caches [9] may be more relevant today than when it was first done, given the increased cost of misses. However, even reducing misses runs up against the problem that increasing miss costs become harder and harder to amortize: if current trends persist, miss costs to DRAM will eventually be in the range of thousands of instructions. If misses become that expensive, it starts to become worth considering options such as taking a context switch on a cache miss. Further, as misses become more expensive, it becomes worth exploring techniques for locking critical data or code in the lowest-level cache.

Taken to their logical conclusion, strategies like software-controlled caches, locking given data or code in the lowest level of cache and context switches on misses imply managing the lowest-level cache as a paged memory.

The approach proposed in this paper, the *RAMpage hierarchy*, is one in which the lowest-level cache (typically made of static RAM, SRAM) is instead considered to be the main memory, and DRAM is considered to be a first-level (L1) paging device. In the RAMpage model, disk

becomes a second-level (L2) paging device.

The RAMpage hierarchy differs from the conventional hierarchy in the way specific levels of the hierarchy are managed; the total amount of hardware is not necessarily significantly different. It is possible that improvements to the basic RAMpage design will require additional hardware, but the same is true of enhancements to a basic cache architecture.

The RAMpage hierarchy will require some modifications to the operating system, since an additional level of paging is introduced. Given the significantly lower miss costs from SRAM to DRAM, as opposed to from DRAM to disk, the optimal page size at the new level will be smaller than traditional page sizes. Also, the faster access time to DRAM as compared with disk implies that trade-offs such as a more complex replacement algorithm versus more misses will differ in the SRAM-DRAM level as compared with the DRAM-disk level.

On the positive side, then, the RAMpage approach offers the potential to address the memory wall without extra hardware cost. The downside is that changes to the operating system are required (i.e., adopting a RAMpage architecture cannot be done simply by changing the hardware). Fortunately, these changes are relatively minor – the paging implementation can follow standard principles with minor differences to take into account the parameters of faults from SRAM to DRAM.

In this paper, a preliminary investigation of the RAMpage hierarchy, in which measurements are made using trace-driven simulation, is presented. Several simplifica-

tions are made to reduce the number of variables in comparing the RAMpage hierarchy to a conventional two-level cache hierarchy. For example, context switches are not accurately modelled, nor are context switches taken on misses. However, a mix of programs is measured, to give an indication of how working sets of a multiprogramming mix develop over time. Also, a trace of a simple page replacement algorithm is used to simulate the costs of replacement in the RAMpage case.

The major aim of the paper is to show that, with the kinds of miss costs that are becoming common today, the RAMpage approach is viable even without pushing it to maximum advantage.

Preliminary findings are promising – taking into account that more attention needs to be paid to detailed simulation. If current trends continue, the RAMpage approach appears to be viable now, even without special hardware support. With extra investment in areas like hardware-supported page tables, the RAMpage approach could be even more attractive.

The remainder of this paper expands on issues raised so far, and presents results of measurement. The next section provides more background on trends, standard memory hierarchies in use today, and the implications of trends on such standard cache-based hierarchies. The following section briefly summarises related work, including software-managed caches and attempts at reducing the impact of the CPU-DRAM gap. After that is a more detailed description of the RAMpage architecture, followed by a section describing the experimental methodology and results. The final section presents conclusions.

## 2 Background

### 2.1 Introduction

Since the mid-1980s, when load-store (RISC) microprocessors became common, the trend in CPU speed improvement has been 50% to 100% per year, while DRAM speed improvement has only been 7% per year [13], leading to a doubling of DRAM reference costs, relative to CPU speed, every 6.2 years [5].

While there have been many attempts at reducing or hiding the costs of DRAM references [26, 6, 18, 5], if current trends continue, it will not be long before miss costs of thousands of instructions are commonplace. For example, it is possible to buy a DEC Alpha system today running at 500MHz with a peak instruction issue rate of 4 instructions per cycle, or an instruction per 0.5ns. At the same time, a typical cache miss cost today is of the order of 200ns. While miss costs are not an exact function of the CPU-DRAM speed gap (improved DRAM organization and bus design for example can reduce miss costs), the trend in the CPU-DRAM speed gap indicates that miss costs of thousands of instructions can be expected in the near future.

The possibility that the growing CPU-DRAM speed gap will eventually result in computer systems running into the memory wall makes it useful to re-evaluate past ap-

proaches which may still be applicable, or which may now be more applicable than when they were first investigated.

The remainder of this section briefly summarizes previous work on software-managed caches, and goes on to consider a range of approaches to reducing the impact of the CPU-DRAM speed gap. The section goes on to present a brief overview of standard cache-based hierarchies with virtual memory, and examines implications of the growing CPU-DRAM speed gap.

### 2.2 Standard Hierarchy

To put the research presented in this paper into perspective, here is a brief overview of common features of current architectures [13]. No attempt is made at covering all possible variations; the focus here is on issues to take into account when comparing RAMpage with a conventional hierarchy.

Typically, systems have one or more levels of cache – most recent designs have two, but three levels are likely to become commonplace as the CPU-DRAM speed gap widens. A cache is organized into blocks (also called lines). A cache reference which hits in the cache essentially divides into three phases:

- *indexing* – determining where in the cache to look
- *tag check* – determining whether the referenced block is present (and possibly updating its status)
- *selection within block* – do the read or write on the portion of the block referenced by the CPU

The CPU issues virtual addresses, which are translated somewhere during the cache reference process. A TLB (translation lookaside buffer) caches page translations, and is used where possible to do page translation within the cache reference cycle. In recent years, architectures have tended to use the virtual address on an index operation, and the physical address in the tags [19], to allow more time for the address translation, though few have gone as far as to use completely virtually addressed caches [9].

There are many other details of caches which could be considered [25, 24]. Of most interest here are replacement strategies in caches, and the role of the TLB.

Caches are generally organized with a varying degree of associativity. No reasonably-sized cache is *fully associative*<sup>1</sup>: the situation where any block can be placed anywhere in the cache. Speed and cost constraints drive cache designers towards a *direct-mapped* cache, where a given block is always found in the same place. A compromise is an *n-way associative cache*. If cache size is  $m$ , and  $1 \leq n \leq m$ , there are  $n$  possible locations for placement of a given block, in an  $n$ -way associative cache. The associativity of a cache impacts on the replacement strategy. If the cache is direct-mapped, replacement is trivial. The higher the associativity, the more useful it is to have a more

<sup>1</sup>A small cache like a TLB is often fully associative.

sophisticated approach to replacement, as there are more choices of blocks *not* to replace.

TLB misses can account for a significant portion of execution time. Some studies have shown that operating system code in particular accounts for a large portion of TLB miss costs [23].

It has also been observed that code which uses large numbers of small data structures randomly scattered over memory can have a high number of TLB misses even if cache misses are low. Such a reference pattern is likely to be common in object-oriented code, especially where no special attention is paid to where objects are placed in memory [7].

In the best case, a TLB miss can be handled from a page table entry in L1 cache; in the worst case, a missing TLB entry may only be present on disk.

### 2.3 Implications of Trends

The growing CPU-DRAM speed gap has various implications for memory system designers.

One important implication is that misses should be reduced as far as possible. Although DRAM references are still much faster than disk references, some of the logic which applies to page faults starts to become an issue for DRAM references. For example, a more expensive replacement strategy in the lowest level of SRAM can be justified if it is more than offset by the reduction in misses. Also, it becomes useful to pin critical resources in the last SRAM level above DRAM.

The way a cache is organized, it is difficult to pin specific blocks into the cache, especially if there is a low degree of associativity. In the extreme case of a direct-mapped cache, pinning anything into the cache means that any other block that maps to the same location can never be present in the cache.

While increasing associativity can address both miss reduction and making it possible to pin blocks in the cache, increasing associativity increases hit time [13]. Furthermore, pinning blocks in a cache would require either an extra tag bit (more hardware) or some other extra overhead on choosing a candidate for replacement.

Concerns of running into the memory wall need to take into account all aspects of the memory hierarchy which may access DRAM – not just the caches but also the TLB.

Page table entries likely to be needed soon but which cannot be accommodated in the TLB are an example of data which could profitably be pinned into the lowest SRAM level. Another example is code and data structures in the operating system for handling context switches.

Ultimately, the only way to deal with the memory wall is to find other work to do on a miss – in other words, to move towards context switches on misses.

Taking all these points into account, there is a strong case for investigating the RAMpage strategy.

### 2.4 Software-Managed Caches

In the 1980s, there was some work on software-managed caches [9]. Software managed caches have the potential to reduce misses by more intelligent replacement than is feasible with hardware-managed caches. As with the RAMpage model, the trade-off is higher replacement costs.

In the 1980s, miss costs were not high enough to make software-managed caches viable. Also, CISC architectures in common use at the time had relatively high costs of traps to the kernel. More recent load-store designs, such as the MIPS architecture, make it possible to do a light-weight trap to the kernel, i.e., a trap does not automatically dump the entire machine state on the stack [19].

The RAMpage model is potentially more viable than the previous approaches to software-managed caches for several reasons. It goes further towards treating the lowest level of SRAM as a fully software-managed level, i.e., a paged memory, with benefits outlined in Section 3. It also benefits from the increase in the CPU-DRAM speed gap since the work on software-managed caches, as well as the potential for relatively low-cost traps in modern CPUs.

### 2.5 Reducing the Impact of the CPU-DRAM Gap

There has been considerable attention directed at reducing the impact of the growing CPU-DRAM speed gap; only a short summary is presented here. Broadly, approaches can be divided into hardware and software strategies.

Hardware strategies can be divided into three major categories; a few examples are given here to illustrate the general principles:

- *reducing miss costs* – variations on DRAM organization including some cache on the DRAMchip, or approaches to capitalize on locality (fast page mode, EDO, etc.); more sophisticated buses; more heavily interleaved RAM
- *reducing number of misses* – more associativity, larger caches
- *hiding miss costs* – victim caches to hold recently replaced data [18]; non-blocking caches [6] to support prefetch or continued out-of-order execution

Software strategies include compiler optimizations [2], algorithm analysis which takes misses into account [21], application restructuring [8], and using page placement to reduce conflict misses [4, 20].

All of these strategies run up against one major fundamental problem: they are at best hiding the underlying trend, not changing it. At some point, a miss to DRAM has to be taken, and when it occurs, if current trends continue, the memory wall remains an issue. The RAMpage approach, with its potential for taking a context switch on a miss, offers the option of a detour around the memory wall.

### 3 The RAMpage Architecture

#### 3.1 Introduction

Given the problems caused by the increasing CPU-DRAM speed gap, it is useful to attempt to find a solution that takes us further than previous approaches.

To summarize the issues raised in Subsection 2.3, it is possible to trade fewer misses for a higher replacement cost, and it would be useful if more operating system traffic could be kept to the SRAM levels of the hierarchy – including TLB misses and context switch code. By making it easier to manage what is at least in the lowest level of SRAM, it should become possible to avoid the worse cases of behaviour of TLB misses (at least in the case where a reference to DRAM or disk is not otherwise required).

Another point to consider is whether it is possible to arrive at a simpler strategy for hits, as compared with a TLB lookup and a cache reference. The result could be faster hits.

Finally, given the growing cost of DRAM references, it may be useful to have the option of taking a context switch on a miss to DRAM. This is not as fanciful as it may at first appear: early virtual memory system's page fault costs were not much different from current miss costs to DRAM<sup>2</sup>.

All these issues lead to the idea of the RAMpage architecture – a revision of the conventional hierarchy in which main memory moves up a level to the lowest level of SRAM, while DRAM becomes the L1 paging device. Disk becomes the L2 paging device.

A major design issue here is to address the problem of high TLB miss costs. In particular, it is a concern that a hit in any SRAM level should not be slowed down by a slow TLB miss. While TLB miss costs may appear to be secondary to the other issues raised, there is little point in improving other aspects of the memory hierarchy, only to have TLB misses remain a bottleneck.

The remainder of this section provides more detail of the RAMpage architecture, followed by some implementation detail. Since this paper aims to investigate feasibility, detail is not developed as fully as if an implementation were currently proposed.

#### 3.2 RAMpage Architecture Fundamentals

It is important to note that a basic RAMpage architecture need not cost more than a conventional cache-based architecture. The major components are the same. It is possible that additional hardware will be justifiable to improve performance, but the same is true of a cache: a simple, basic design can be enhanced by adding extra hardware, such as a victim cache to hold recently replaced blocks [18]. In fact, the equivalent concept to a victim cache is commonly implemented in paged systems without extra hardware (re-

<sup>2</sup>Astonishing though it may seem, in the 1950s, drum memories were being produced with sub-millisecond access times: see for example <<http://www.cc.gatech.edu/gvu/people/Faculty/Randy.Carpenter/folklore/dec92-vln2.html>>.

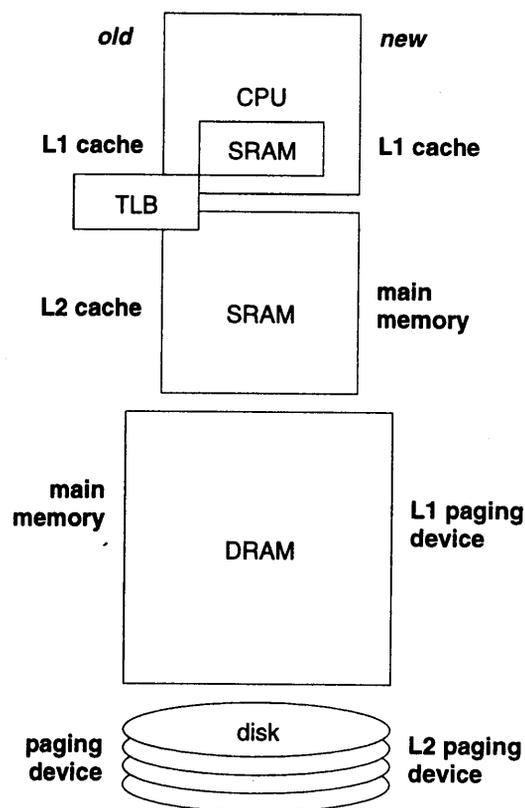


Figure 1: A Standard Hierarchy vs. a RAMpage Hierarchy

cently replaced pages are kept on a *standby page list*; the page which was on the list longest is the one actually discarded [10]).

Figure 1 illustrates the traditional versus the RAMpage hierarchy. The major components (as illustrated here) are the same; detailed implementation differs.

To expand on Figure 1, the major difference at the hardware level is that the RAMpage architecture does not use cache tags in the lowest level of SRAM. However, whether there is a net increase or decrease in SRAM used for page tables as opposed to tags depends on the page size of the RAMpage system, an issue which is taken up in Section 4, where results are presented. The approach here is to use the same amount of SRAM in all simulations, to avoid having cost as a major issue. A potential win though is that in recent designs, the L2 tags are on-chip [16, 15], whereas the RAMpage architecture does not require that page tables be stored on-chip (except those entries stored in the TLB which is present in any case in the traditional architecture).

At the software level, the major difference between RAMpage and conventional architectures is that an additional level of page table is needed.

Since the physical address space of the SRAM main memory is relatively small compared with the total address space, a sparse page table approach is needed. In this work, an inverted page table [14] is used to map the entire SRAM main memory to virtual addresses.

It is useful to ensure that the part of the page table needed for any hit to any SRAM level is in the SRAM

main memory. Even if the inverted page table needs to do more work for a lookup than a forward page table, the worse cases of fetching a page table entry from DRAM or from disk will never occur for a reference which hits in one of the SRAM levels.

In the RAMpage model, a hit in any SRAM level, in the best case (in terms of TLB behaviour), proceeds from a TLB hit to a reference in either one of the the SRAM levels memory. In the worst case, a hit in an SRAM level requires that a TLB miss be handled from the lowest level of SRAM.

While a direct-mapped cache may be able to match the RAMpage architecture on hit cost, we argue that our organization is simpler and therefore easier to make fast in the best case, and the worst case is also better.

Once associativity is introduced into the cache, it becomes even harder to achieve a fast hit time. By comparison, the RAMpage approach is fully associative, which should reduce misses, yet is simpler in the case of a hit than even a direct-mapped cache. The cost is that replacements are more expensive than with a conventional cache architecture, as replacements are handled in software – in the style of a conventional paging implementation.

The RAMpage approach is slower on a TLB miss compared with a conventional architecture's best-case TLB miss, where the required page table entry is cached. In this case, the slower lookup of the RAMpage inverted page table results in lower performance for the RAMpage hierarchy. The number of TLB misses will generally be higher in a RAMpage architecture than a conventional architecture unless the RAMpage SRAM pages are at least as big as traditional pages.

The other areas where RAMpage will take a performance hit are in handling replacements and misses. Replacements are slower than with a cache since a software approach is used. Misses are slower since a page needs to be larger than typical cache blocks, as is illustrated in Section 4.

### 3.3 Preliminary Implementation Detail

Cache levels above the SRAM main memory are organized the same way in the RAMpage model as in conventional memory hierarchies.

The lowest SRAM level is organized like a conventional DRAM paged memory, except some changes are made to take into account the fact that the SRAM level is relatively small as compared to DRAM, and misses to DRAM are relatively quick as compared to page faults to disk. Also, some attention is paid to reducing TLB miss costs, especially where the reference concerned does not require access to DRAM or disk.

An inverted page table is used to implement page translation as, in this way, it can be guaranteed that any page present in the SRAM main memory is mapped in the portion of the page table present in SRAM main memory. A conventional forward-mapped table for a hierarchy where the main memory is relatively small would in any

case have to be modified to take into account the relatively small portion of the total virtual address space resident in the main memory.

Compared with a cache, an SRAM main memory does not require a cache tag index operation, checking for a hit, or extracting the appropriate sub-block. Once a TLB translation has been performed, the appropriate reference can be performed directly. While cache indexing and TLB lookup can be performed concurrently in a virtually indexed cache, checking for a hit and subblock extraction are overheads in the conventional model which are avoided in the RAMpage model. For simplicity, the simulation implemented for this work does not model the speed difference between RAMpage and cache hits. However, the greater simplicity of the RAMpage model should aid the hardware designer in improving performance of hits.

RAMpage replacements are handled using a simple text book clock page replacement algorithm [10, 22]. The philosophy is that a simple approach with minimal cost, at the expense of slightly more misses, is the right approach while miss costs are still relatively low. As miss costs to DRAM increase, more sophisticated replacement strategies will be justified.

For purposes of this work, the L2 paging device – the disk – is not implemented, as this level of the hierarchy does not differ from that of a conventional cache-based architecture.

It is possible – depending on the page size – that an additional level of page table may need more memory than cache tags. The trade offs are further explored in Section 4. It should be noted however that page table entries would be present anyway in the lowest level of cache, so the trade-off is hard to determine in general. In the conventional cache-based system, these page table entries would only be for paging in the DRAM and disk levels, whereas in the RAMpage model, page table entries for both levels of paging might be present in the SRAM main memory. However, most page translations in the RAMpage model are likely to be for pages resident in the SRAM main memory; the L2 page translations are not needed for example to service TLB misses, except in the case where there is also a miss from SRAM to DRAM or disk.

## 4 Experimental Evaluation

### 4.1 Introduction

Since there are many variables in a memory hierarchy which could confuse evaluation of the RAMpage alternative, the RAMpage implementation and a cache-based hierarchy used for comparison have been kept as simple as possible.

To simplify variation in the hierarchy, trace-based simulation is used. Traces are drawn from SPEC benchmarks<sup>3</sup>. Each trace is approximately 1-million references. A to-

<sup>3</sup>A selection of 20 traces from SPEC benchmarks produced by Mark Hill was used, from the FTP location <ftp://tracebase.nmsu.edu/pub/traces/uni/r3000/pdt/>.

tal of 20 traces is used, with a mix of floating point and integer code. 1-million traces is not much to simulate a realistic working set, so we produced longer traces with an approximation to trace stitching [1], by using the same trace repeatedly. Since this technique is obviously of limited accuracy, we do not rely heavily on these longer traces to draw conclusions.

Traces from each benchmark are interleaved, with a fixed number of references from each before switching to another, to simulate a multiprogramming workload. We have made no attempt at simulating the cost of context switches (other than the impact on contents of each level of the memory hierarchy).

The remainder of this section details specifics of the architectures as implemented for measurement, followed by measurement and discussion of the results.

## 4.2 Simulated Architectures

The simulated CPU cycle time is 2ns (500MHz). For simplicity, we have not simulated a superscalar architecture, nor have we attempted to take into account pipeline interactions with the memory system (other than stalls for misses).

The TLB is identical in both memory hierarchies. It is software managed, fully associative and contains 64 virtual to physical address translations which are shared between code and data references. TLB hits require a single processor cycle, and misses are managed by software which retrieves the required translation from the page table. The TLB in the conventional hierarchy maps DRAM pages, whereas the RAMpage TLB maps SRAM pages.

Both hierarchies employ the same primary cache organization. The on-chip cache is physically addressed, split (16Kbytes each for the I- and D-caches) and direct mapped. A write back strategy with a write allocate miss policy is adopted. Block size is fixed at 32bytes, with a full block fetched on a miss. Dirty misses cause the entire block to be written back to the next level. Both the I- and D-caches cycle at the same rate as the CPU, with read hits taking one cycle and write hits requiring two. Misses incur processor stalls until the required block has been loaded.

The L2 cache in the standard hierarchy is external, physically addressed and unified. It is direct mapped and write back with a write allocate miss policy. Cache size is fixed at 4 Mbytes with block size varied from 128bytes. Like the primary cache, full blocks are fetched on a miss and dirty misses result in the entire block being written back to DRAM. Cycle time is set at 30ns, and so 1 secondary cache cycle = 15 CPU cycles. Read hits require a single secondary cache cycle, and write hits twice that amount. An inclusion property between L1 and L2 caches is maintained (i.e., the L1 cache is a subset of the L2 cache, other than that some L1 blocks may be dirty with respect to L2).

Main memory in the RAMpage model has the same *total capacity* as the secondary cache in the baseline system with 128byte blocks; this includes the space taken up

by the L2 tags. It is assumed that each tag in the L2 cache requires 4bytes of storage, and so the 32K frames (each 128bytes) require a total of  $32K \times 4\text{bytes} = 128K\text{bytes}$  of storage. Hence, the total capacity of the secondary cache is 4.125 Mbytes. Page size is varied in powers of two from 128bytes to 4096bytes. Inclusion is also maintained between the L2 cache and the SRAM main memory.

Adjacent layers of the hierarchies are connected by a 30ns bus which is 32bytes wide. Thus, transfer of a cache block between the L1 and the L2 cache (or SRAM main memory) requires a single bus cycle; transferring a block between L2 (or SRAM main memory) and DRAM requires four or more bus cycles (depending on block size).

An infinite DRAM is simulated, in which all data and code is preloaded, to avoid the need to simulate accesses to disk.

DRAM latencies are varied at 50 cycle increments from 65 cycles for reads and 60 cycles for writes, to 315 cycles for reads and 310 cycles for writes.

## 4.3 Results

In this subsection, results of trace-driven simulation are presented. These results are intended to provide a first cut at evaluating the RAMpage hierarchy, by comparison with a simple cache-based hierarchy.

Results are presented with a range of block (line or page) sizes in the lowest-level SRAM (L2-cache, or main memory, depending on the model being evaluated).

Table 1 contains simulated run times of both a conventional cache hierarchy and a RAMpage hierarchy. In each case block size at the L2 (or SRAM main memory) level is varied, and memory speed is varied. Run times are given in simulation time, rather than wall clock units.

It is interesting to observe that simulated times for the best block (or page) size for both hierarchies differ generally by about 1%. However, the block size for which the best performance occurs differs markedly between the RAMpage and cache-based architectures. The conventional architecture achieves its best performance for the smallest block size used here, whereas RAMpage achieves the best performance for a 512byte page for slower miss costs, and a 1024 byte page for faster miss costs.

These numbers are to some extent an artifact of the simulation: no account has been taken of memory system enhancements which reduce the cost of subsequent references to DRAM, after the initial reference of a sequence of contiguous references. While all cases will be improved by a more realistic DRAM simulation, the RAMpage architecture will be improved more since it performs better on larger blocks (pages).

Table 2 illustrates how increasing the L2 block size or SRAM page size impacts on number of DRAM references. Of particular interest is that for the relatively small traces used here, there are no write backs at all in the RAMpage case, for smaller page sizes. Only with page sizes of 2K or more are there write backs, which confirms that the clock algorithm is more efficient in selecting appropriate candi-

DRAM Speed	128bytes	256bytes	512bytes	1024bytes	2048bytes	4096bytes
rd: 65; wr: 60	71 101 033 85 858 027	72 079 551 76 754 297	74 144 325 73 294 034	77 871 338 71 524 452	85 153 856 73 991 456	100 005 393 79 562 784
rd: 115; wr: 110	77 159 433 90 714 428	78 919 151 82 032 697	82 329 925 79 154 034	88 574 382 78 260 452	100 305 856 83 565 856	124 523 793 92 253 984
rd: 165; wr: 160	83 262 569 95 582 228	85 823 916 87 311 097	90 454 725 85 014 034	99 093 738 84 996 452	115 457 856 93 140 256	148 830 528 104 945 184
rd: 215; wr: 210	89 311 769 100 450 028	92 663 916 92 589 497	98 641 125 90 874 034	109 803 182 91 732 452	130 609 856 102 714 656	173 437 524 117 636 384
rd: 265; wr: 260	95 298 633 105 317 828	99 503 916 97 867 897	106 827 525 96 734 034	120 417 582 98 468 452	145 761 856 112 289 056	197 911 124 130 327 584
rd: 315; wr: 310	101 410 169 110 185 628	106 343 916 103 146 297	115 072 325 102 594 034	131 031 982 105 204 452	160 913 856 121 863 456	222 384 724 143 018 784

Table 1: Hierarchy Performance. Total processor cycles after 20 million references, with varying DRAM speed and L2 page/block size, and context switches every 100000 references. Values for the standard hierarchy appear in the upper portion of each row, and those for the new hierarchy in the lower.

L2 page/block size	Standard Hierarchy		RAMpage Hierarchy	
	Reads	Writes	Reads	Writes
128bytes	106 052	14 396	97 356	0
256bytes	118 256	18 536	105 568	0
512bytes	138 640	25 072	117 200	0
1024bytes	173 920	38 304	133 984	0
2048bytes	238 592	64 448	168 128	23 360
4096bytes	369 536	120 832	204 416	49 408

Table 2: Total DRAM references (counted as number of bus transactions) incurred by the two hierarchies after simulation of 20-million references.

dates for write backs than is a simple direct-mapped cache.

The number of references here is not a direct indicator of performance, as even with our inefficient bus, multiple references forming a single miss or write back take less time than the equivalent number of single references. Thus, the larger block or page sizes gain as long as increased number of references per miss is offset by the greater efficiency of handling multiple references for a given transaction.

A more significant factor as page size increases is the reduction in TLB misses, which is the largest factor in the improved performance of the RAMpage model for larger page sizes. This variation does not occur with the conventional architecture, since its TLB maps DRAM pages, the size of which is fixed.

As can be seen in Figure 2, TLB hits increase significantly in the RAMpage case, as page size increases. However, TLB hits remain constant on the standard hierarchy as block size changes (since DRAM pages in this case are mapped in the TLB, and this size is held constant).

It should be noted that the TLB miss rates are not directly comparable across the two architectures. Of more interest is the actual TLB miss cost, since the RAMpage architecture is designed to accommodate all page translations within SRAM (provided the page concerned is in SRAM). Figure 3 illustrates how the new hierarchy performs badly with small pages. With small pages, there are many TLB misses. The data as given here shows TLB misses in terms of extra references generated by the TLB miss handler. A

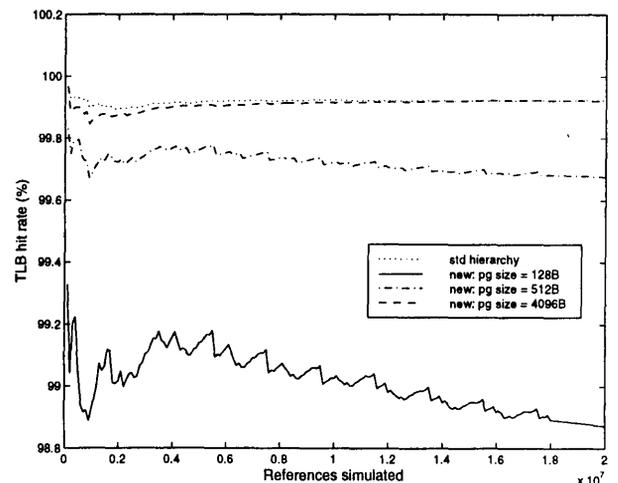


Figure 2: TLB Hit Rate. The TLB behaviour is the same for all block sizes of the conventional hierarchy.

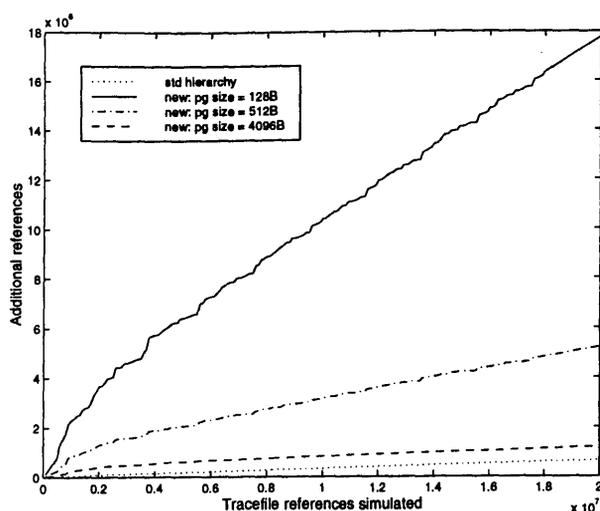


Figure 3: Cumulative TLB Miss Cost (in terms of extra instruction and data references). The TLB behaviour is the same for all block sizes of the conventional hierarchy.

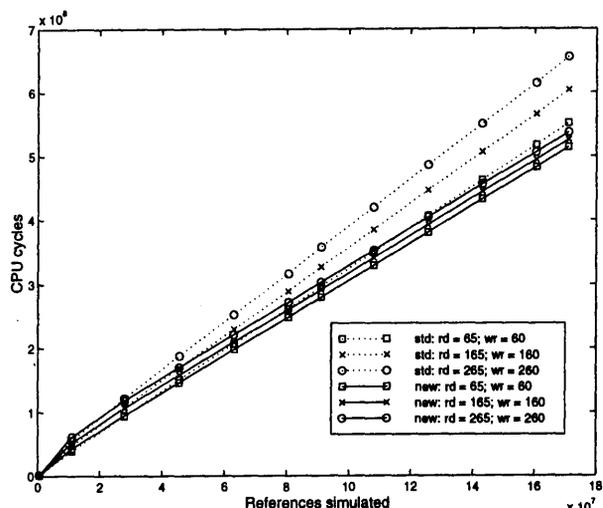


Figure 4: Cumulative Simulated Run Time (20-million references).

smaller fraction of the extra references for the RAMpage case go to DRAM, so the higher number of extra references generated by the TLB miss handler for page sizes of 512bytes and above is not significant.

Finally, it is of interest to see how performance develops for a longer trace.

Figure 4 illustrates how the new hierarchy starts to become a win once cold start misses are no longer an issue. This matches our view that the RAMpage architecture should come out ahead without too much work to optimize it. The data here is for the best cases of each hierarchy: 128byte blocks for the L2 cache, and 512bytes for the RAMpage model.

Two details are of particular interest. Even though the miss cost is increased by a factor of more than 4, the RAMpage hierarchy only slows down by about 5%. For the highest miss cost, the RAMpage hierarchy is about 25%

faster than the conventional hierarchy. The overall speed improvement results from a combination of trade-offs:

- each miss costs more in RAMpage, but there are fewer misses, especially as SRAM page size grows
- there are more TLB misses in RAMpage, especially at small SRAM page sizes, which also favours larger pages – but TLB misses on average cost less than in the conventional architecture

#### 4.4 Discussion of Results

Despite the simplifications made to limit the variables in this initial work, the data presented here provides evidence that the RAMpage hierarchy is viable, especially as miss costs increase. If the stitched traces are typical of real runs, the simplified RAMpage architecture simulated here comes out ahead, even with miss costs seen on some current machines.

While TLB misses are a significant overhead in the new architecture, the simulation results show that, provided the page size is large enough, the extra number of TLB misses is not high enough to offset the gain in having most TLB misses serviced from SRAM. While large page sizes may cause extra overhead on misses, they also save overhead on page table size; fine-tuning details of this kind will require more detailed simulation, including a more realistic model of the memory system.

### 5 Conclusions

#### 5.1 Introduction

The RAMpage work presented here is a beginning. Initial results indicate that it is worth further researching the RAMpage approach. This concluding section summarizes the major issues and results contained in this paper, then goes on to outline planned future work. In conclusion, the results are put into context.

#### 5.2 Summary

The RAMpage memory hierarchy offers a way around the memory wall, by making it possible to do context switches on misses. This paper has shown that even without taking context switches on misses, the underlying infrastructure of the RAMpage approach has competitive performance with a simple conventional L2 cache-based system. Both the cache implementation and the RAMpage implementation could have many improvements made to them to make them more realistic, but the basic principle demonstrated in this paper is that the RAMpage approach is feasible.

#### 5.3 Future Work

An important issue not addressed in this paper is that of context switches on misses. If miss costs are high enough (in this case, page faults from SRAM main memory to

DRAM), it becomes worth while to take a context switch on a miss. Work is in progress to investigate the conditions in which it becomes worth taking a context switch on a miss to DRAM. Additions to the current simulation include more accurate simulation of context switch code and operating system data structures.

Another project under consideration is hardware support for inverted page tables. In this instance, the likely focus of the project will be on the use of formal methods to ensure correct implementation, since the design is likely to be relatively straightforward.

A possible benefit of the RAMpage approach is that object-oriented code will have fewer expensive TLB misses, given the RAMpage goal of ensuring that all pages currently found at least in part in any SRAM level have their translation at worst in the SRAM main memory. Users of languages such as Smalltalk or Java which use automatic memory allocation will particularly benefit from this improvement. Our future work will include investigation of the TLB behaviour of object-oriented code on a RAMpage versus a conventional cached machine.

Multiprocessor systems introduce another range of problems. To minimize variables in this initial research, multiprocessor systems are not considered. However, some lessons could be drawn from work on distributed shared memory (DSM) systems [3, 11].

#### 5.4 Final Summary

Overall, the RAMpage approach appears to be promising. Further work is needed to investigate missing details. However, given initial data presented in this paper, there is reason to continue with further investigation.

The fact that an improvement of up to 25% was seen over the simple cache architecture is an indication that the RAMpage approach is viable, even without going to the extent of implementing currently missing details like context switches on misses. More important, the fact that RAMpage scales up better than the conventional cache-based architecture as miss costs grow, indicates that the RAMpage strategy will become more viable if current trends persist.

It somehow seems a more satisfactory metaphor to walk around a wall than to run into it.

## References

- [1] A. Agarwal, J. Hennessy, and M. Horowitz. Cache performance of operating system and multiprogramming workloads. *ACM Transactions on Computer Systems*, 6(4):393–431, November 1988.
- [2] D.F. Bacon, S.L. Graham, and O.J. Sharp. Compiler transformations for high performance computing. *Computing Surveys*, 26(4):345–420, December 1994.
- [3] J.K. Bennet, J.B. Carter, and W. Zwaenepoel. Adaptive software cache management for distributed shared memory architectures. In *Proceedings of the 17th International Symposium on Computer Architecture*, pages 125–134, Seattle, WA, May 1990.
- [4] B.N. Bershad, D. Lee, T.H. Romer, and J.B. Chen. Avoiding conflict misses dynamically in large direct-mapped caches. In *Proceedings of the 6th International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 158–170, October 1994.
- [5] K. Boland and A. Dollas. Predicting and precluding problems with memory latency. *IEEE Micro*, 14(4):59–67, August 1994.
- [6] T. Chen and J. Baer. Reducing memory latency via non-blocking and prefetching caches. In *Proceedings of the 5th International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 51–61, September 1992.
- [7] D.R. Cheriton, H.A. Goosen, H. Holbrook, and P. Machanick. Restructuring a parallel simulation to improve cache behavior in a shared-memory multiprocessor: The value of distributed synchronization. In *Proceedings of the 7th Workshop on Parallel and Distributed Simulation*, pages 159–162, San Diego, May 1993.
- [8] D.R. Cheriton, H.A. Goosen, and P. Machanick. Restructuring a parallel simulation to improve cache behavior in a shared-memory multiprocessor: A first experience. In *Proceedings of the International Symposium on Shared Memory Multiprocessing*, pages 109–118, Tokyo, April 1991.
- [9] D.R. Cheriton, G. Slavenburg, and P. Boyle. Software-controlled caches in the VMP multiprocessor. In *Proceedings of the 13th International Symposium on Computer Architecture*, pages 366–374, Tokyo, June 1986.
- [10] C. Crowley. *Operating Systems: A Design-Oriented Approach*. Irwin Publishing, 1997.
- [11] S. Dwarkadas, P. Keleher, A.L. Cox, and W. Zwaenepoel. Release consistent software distributed shared memory on emerging network technology. In *Proceedings of the 20th International Symposium on Computer Architecture*, pages 144–155, San Diego, CA, May 1993.
- [12] J.L. Hennessy and N.P. Jouppi. Computer technology and architecture: An evolving interaction. *Computer*, 24(9):18–29, September 1991.
- [13] J.L. Hennessy and D.A. Patterson. *Computer Architecture: A Quantitative Approach*. Morgan Kaufmann Publishers, San Mateo, CA, 2nd edition, 1996.

- [14] J. Huck and J. Hays. Architectural support for translation table management in large address space machines. In *Proceedings of the 20th International Symposium on Computer Architecture*, pages 39–50, May 1993.
- [15] IBM. PowerPC 740 and PowerPC 750 SCM RISC microprocessor: Hardware specifications. [http://www.chips.ibm.com/products/ppc/documents/datasheets/750/740\\_750\\_hwspec\\_rev10.gen.pdf](http://www.chips.ibm.com/products/ppc/documents/datasheets/750/740_750_hwspec_rev10.gen.pdf), January 1998.
- [16] Intel. Pentium II processor product overview. <http://developer.intel.com/design/PentiumII/prodbref/index.htm>, 1998.
- [17] E.E. Johnson. Graffiti on the memory wall. *Computer Architecture News*, 23(4):7–8, September 1995.
- [18] N.P. Jouppi. Improving direct-mapped cache performance by the addition of a small fully-associative cache and prefetch buffers. In *Proceedings of the 17th International Symposium on Computer Architecture*, pages 364–373, May 1990.
- [19] G. Kane and J. Heinrich. *MIPS RISC Architecture*. Prentice Hall, Englewood Cliffs, NJ, 1992.
- [20] R.E. Kessler and M.D. Hill. Page placement algorithms for large real-indexed caches. *ACM Transactions on Computer Systems*, 10(4):338–359, November 1992.
- [21] M.S. Lam, E.E. Rothberg, and M.E. Wolf. The cache performance and optimizations of blocked algorithms. In *Proceedings of the 4th International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 63–74, Santa Clara, CA, 1991.
- [22] M. Maekawa, A.E. Oldehoeft, and R.R. Oldehoeft. *Operating Systems: Advanced Concepts*. Benjamin/Cummings Publishing, Menlo Park, CA, 1987.
- [23] D. Nagle, R. Uhlig, T. Stanley, S. Sechrest, T. Mudge, and R. Brown. Design tradeoffs for software-managed TLBs. In *Proceedings of the 20th International Symposium on Computer Architecture*, pages 27–38, May 1993.
- [24] S. Przybylski, M. Horowitz, and J. Hennessy. Design tradeoffs in cache design. In *Proceedings of the 15th International Symposium on Computer Architecture*, pages 290–298, May/June 1988.
- [25] A.J. Smith. Cache memories. *Computing Surveys*, 14(3):473–530, September 1982.
- [26] G.S. Sohi and M. Franklin. High-bandwidth data memory systems for superscalar processors. In *Proceedings of the 4th International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 53–62, April 1991.
- [27] M.V. Wilkes. The memory wall and the CMOS end-point. *Computer Architecture News*, 23(4):4–6, September 1995.
- [28] W.A. Wulf and S.A. McKee. Hitting the memory wall: Implications of the obvious. *Computer Architecture News*, 23(1):20–24, March 1995.

Received: 10/97, Accepted: 3/98

## Acknowledgements

Ian Sanders gave useful comments on this paper. An anonymous referee also gave some useful input.

## Notes for Contributors

---

The prime purpose of the journal is to publish original research papers in the fields of Computer Science and Information Systems, as well as shorter technical research notes. However, non-refereed review and exploratory articles of interest to the journal's readers will be considered for publication under sections marked as Communications of Viewpoints. While English is the preferred language of the journal, papers in Afrikaans will also be accepted. Typed manuscripts for review should be submitted in triplicate to the editor.

### Form of Manuscript

Manuscripts for *review* should be prepared according to the following guidelines:

- Use wide margins and 1½ or double spacing.
- The first page should include:
  - the title (as brief as possible)
  - the author's initials and surname
  - the author's affiliation and address
  - an abstract of less than 200 words
  - an appropriate keyword list
  - a list of relevant Computing Review Categories
- Tables and figures should be numbered and titled.
- References should be listed at the end of the text in alphabetic order of the (first) author's surname, and should be cited in the text according to the Harvard. References should also be according to the Harvard method.

Manuscripts accepted for publication should comply with guidelines as set out on the SACJ web page,

<http://www.cs.up.ac.za/sacj>

which gives a number of examples.

SACJ is produced using the L<sup>A</sup>T<sub>E</sub>X document preparation system, in particular L<sup>A</sup>T<sub>E</sub>X 2<sub>ε</sub>. Previous versions were produced using a style file for a much older version

of L<sup>A</sup>T<sub>E</sub>X, which is no longer supported. Please see the web site for further information on how to produce manuscripts which have been accepted for publication.

Authors of accepted publications will be required to sign a copyright transfer form.

### Charges

Charges per final page will be levied on papers accepted for publication. They will be scaled to reflect typesetting, reproduction and other costs. Currently, the minimum rate is R30.00 per final page for contributions which require no further attention. The maximum is R120.00, prices inclusive of VAT.

These charges may be waived upon request of the author and the discretion of the editor.

### Proofs

Proofs of accepted papers may be sent to the author to ensure that typesetting is correct, and not for addition of new material or major amendments to the text. Corrected proofs should be returned to the production editor within three days.

### Letters and Communications

Letters to the editor are welcomed. They should be signed, and should be limited to about 500 words. Announcements and communications of interest to the readership will be considered for publication in a separate section of the journal. Communications may also reflect minor research contributions. However, such communications will not be refereed and will not be deemed as fully-fledged publications for state subsidy purposes.

### Book Reviews

Contributions in this regard will be welcomed. Views and opinions expressed in such reviews should, however, be regarded as those of the reviewer alone.

### Advertisement

Placement of advertisements at R1000.00 per full page per issue and R500.00 per half page per issue will be considered. These charges exclude specialised production costs, which will be borne by the advertiser. Enquiries should be directed to the editor.

---

## Contents

### Editorial

<b>N. du Plooy</b> .....	1
--------------------------	---

---

### Research Articles

Text Categorization as an Information Retrieval Task <b>H. Pajmans</b> .....	4
---	---

Preliminary Investigation of the RAMpage Memory Hierarchy <b>P. Machanick and P. Salverda</b> .....	16
--	----

Changes in Entry-Level University Students' Attitudes to Computers from 1985 to 1997 <b>M.C. Clarke and G.R. Finnie</b> .....	26
---	----

A Performance Analyser for the Numerical Solution of General Markov Chains <b>W. Konttenbelt and P. Kritzinger</b> .....	34
---	----

Specific Acquisition of Collective Belief Knowledge for Socially Motivated Multiagent Systems <b>V. Ram</b> .....	44
---	----

---

### Communications and Viewpoints

Fractal Image Compression <b>E. Cloete and L.M. Venter</b> .....	A49
---	-----

Applied Lambda Calculus: Using a Type Theory Based Proof Assistant <b>L. Pretorius</b> .....	A55
---	-----

Recursive Specifications and Formal Logic: What Benefits for Intelligent Tutoring Systems? <b>Lot Tcheeko</b> .....	A63
---	-----

---