



A STUDY OF THE EFFECT OF USER PROGRAM OPTIMIZATION IN A PAGING SYSTEM

Lesin W. Comeau
International Business Machines Corporation
Cambridge, Massachusetts

Summary

Much attention has been directed to paging algorithms and little to the role of the user in this environment. This paper describes an experiment which is an attempt to determine the significance of efforts by the user to improve the paging characteristics of his program.

The problem of throughput in a computing system is primarily one of balancing the flow of data and programs through a hierarchy of storages. The problem is considered solved when for every available processor cycle there is a matching demand for that cycle in the primary (execution) store. Since programs and their data usually originate in a location other than the execution store, there is a delay associated with the movement of data and programs to the primary store. The delay has two components, the operational speed (data transfer time) and the positioning, or access time, of the secondary storage device. Since the access time usually exceeds the data transfer time by an order of magnitude, the problem of transferring information to the primary store has been named the "access gap" problem.

Previous software solutions to the access gap problem divided the problem into two areas of responsibility. First, the problem programmer was responsible for buffering and blocking his data such that high CPU utilization was realized. Subsequent studies showed that this alone would not result in a significant overall increase in system throughput. As technology improved and CPU's and execution store cycles were speeded up, it became apparent that job set-up and initialization was responsible for low throughput performance. A second solution to the "access gap" problem, multiprogramming, was therefore implemented. In this approach, the system programmer is given the responsibility of buffering programs in much the same way the user (problem programmer) buffers data.

It is evident that the critical parameter in this multiprogramming environment is that of block size, since this will determine the number of buffers which can be created from available memory, which is the same as the level of multiprogramming present (i. e. the number of programs which can simultaneously reside in primary storage). The usual blocking factor in a multiprogramming environment is 'one'; that is, the program is treated as one block and a core buffer space is reserved for this amount. In much the same manner that the problem programmer finds when he blocks data this way, we find that programs which result in large block size usually have low core utilization with a resultant reduction in throughput. The low core utilization comes about in two ways. First, a portion of the program is normally never executed, and hence its primary store occupancy is not required. Second, sections of the program are linear in execution with respect to time, and once they have been used, they are not required again; but in systems in which programs are blocked in one unit this unnecessary code must remain in the primary store.

A method of artificial blocking has been devised in an attempt to solve this blocking factor problem incurred with programs. This technique, paging, first used to extend the address space of a small physical core¹, is now used to reduce the physical core requirements of programs written in a very large address space. The alternative to this approach is for the programmer to utilize only a small portion of the address space available (with resultant increase in the core available for multiprogramming) and to substitute an overlay structure for paging. This is distasteful for two reasons. First, it requires a pre-planning of routines -- which to overlay and which to leave resident. Pre-planning in large programming systems is usually inadequate. Second, if only a small part of the available address space is utilized and there are no jobs to be multiprogrammed, then the remaining physical core is not utilized properly. Solution time for the problem is thereby needlessly increased.

There is in a paging system, however, one gross source of error. It is found in the very premise which makes paging in a large address space attractive in the first place, namely, the block size. Ideally, a multiprogrammed paging system attempts to pick a page size which closely emulates a natural grouping of instructions. Since the natural grouping is usually found to be quite small², a compromise page or block size must be derived, as a solution; otherwise the mapping hardware or the storage requirements for mapping tables involve considerable expense.

Unfortunately, those who emphasize paging as a solution to the multiprogramming problem fail to define the responsibility of the user in this environment.

It has been stated that the programmer can do nothing easily to improve the situation.³ It is the contention of the author that this is not so and that, in fact, it is the responsibility of a programmer who executes in an algorithmically managed storage to adjust his programs to the page size if they are to be run frequently.

To prove this contention, the Cambridge Monitor System (CMS)⁴ which was written to take no advantage of a paging environment, was picked for experimentation because its authors were resident and because it was the conversational monitor which would receive highest use on our system. We first asked one of the knowledgeable systems programmers to arrange the nucleus such that the resulting load list would place routines dependent on each other in the same 4k byte block, the page size of the 'virtual machine' system of the Cambridge Scientific Center^{5,6}. The result of this, the aligned deck, is shown in Figure 1. Subsequently a data reduction program PSTA (Appendix A) was written which produces a map of core (see Fig. 2) as paging demands are incurred in program execution. In order to produce these demands, the amount of physical core available to the program was restricted to 64K (16 pages).

A second CMS programmer was given the PSTA output and told to structure his load list (i. e., to rearrange the load sequence of subroutines) so as to reduce the total page requirement. In three passes of data reduction, (PSTA output) a 50% reduction in the number of page transfers was achieved in this JAS arrangement from the aligned version (Fig. 1).

The experiment involved the assembly of a single deck in all runs. Only the nucleus which consisted of 64 movable subroutines, was resequenced. The assembler, which runs in conjunction with the nucleus, was not changed.

The total machine time involved was approximately 20 minutes, and the number of man hours totaled approximately 24, for the JAS deck. The aligned deck required only .5 man hours to arrange.

Since there had been little work in structuring programs to run in a paging environment, this effort was considered worthy of further investigation, namely, into the effect of applying no effort to achieve optimization. Since there were 64! possible arrangements with CMS nucleus, it is impossible to determine, if a deck was arranged in a random manner, the probability of doing better than the JAS arrangement; the following two experiments did show, however, that it could be significantly worse. The deck was sequenced in a random manner, and in an alphabetic fashion (which is a common arrangement that facilitates deck replacement while debugging); the results, as compared to the easily achieved improvement are rather significant (Fig. 1).

To restate the conditions of the experiment shown in Figure 1:

- | | |
|-------------|---|
| Alphabetic: | the decks were arranged in the alphabetic sequences of the names punched in columns 73-80. |
| Random: | decks were given sequential numbers from the JAS arrangement, then resequenced by a random number table. |
| Aligned: | the programmer rearranged his decks based on a knowledge of 4K byte pages. |
| JAS: | programmer, given a map of paging actions generated while his job was in execution, resequenced his load list for an improved paging performance. |

The same experiment, using a Fortran job, was run to measure the difference between alphabetic and JAS.

Although the difference of 3.1 in this case is less significant, it must be noted that a level of optimization which was achieved in one environment (i. e., the assembly) had a beneficial effect in the other. This is what one would expect since the section of code optimized consisted of the interrupt and data management routines used by both the assembler and Fortran.

In view of such significant results, it must be emphasized that no code has been rewritten; only the relative order of the subroutines within a deck has been changed.

It is, therefore, the author's opinion that the user optimization is not only easily achieved, but absolutely necessary for frequent operation in a paging environment. Effort is now being directed to constructing an algorithm and program to perform the load list restructuring mechanically.

The author is most grateful to the CMS people who participated in this experiment, especially to Mr. J. B. Harmon, the project manager, who provided the able assistance of Mr. J. A. Seymour, who structured the JAS deck.

References

1. T. Kilburn, D. B. G. Edwards, M. J. Lenigan, and F. H. Sumner, "One Level Storage System" - IRE Transactions on Electronic Computers, Vol. EC-12, PP. 223-235, April, 1962.
2. O'Neill, R. W. - "Experience Using A Timesharing Multiprogramming System with Dynamic Address Relocation Hardware". AFIPS Conference Proceedings, Vol. 30, PP. 611-621, April, 1957.
3. G. H. Fine, C. W. Jackson, and P. V. McIssac, "Dynamic Program Behavior Under Paging", System Development Corporation SP-2397, June, 1966.
4. Cambridge Monitor System - A user's manual available from the Cambridge Scientific Center.
5. A. B. Lindquist, R. R. Seeber, and L. W. Comeau - "A Time-Sharing System Using an Associative Memory", Proceedings of the IEE, Special Issue on Computers, Vol. 54, No. 12, PP. 1774-1779, December, 1966.
6. R. J. Adair, R. U. Bayles, L. W. Comeau, and R. J. Creasy - "A Virtual Machine System for the 360/40", Cambridge Scientific Center Report 36.010, May, 1966.




FIGURE I.

APPENDIX A
PAGING STATISTICS PROGRAM

- PSTA -

This program takes the output tape created by CP/40⁵ which is a record of significant information present when a demand for a new page is incurred in execution by the "virtual machine." The output is printed one line per page demand in the following manner:

INPAG	the page requested by the virtual machine
OUTPAG	the page picked by the algorithm to be overlaid by the incoming page
INSTR	the instruction counter of the "virtual machine" when the page demand occurs
Page Map	This is a record of the actual pages in the core storage when the demand occurred. The swap is already indicated. Page numbers proceed from left to right with each asterisk representing a page (Page numbers 0-3F). A plus indicates the used bit is one, a zero indicates used bit equals zero, and a space (blank) indicates that page is not now resident in core storage.

CORE RESIDENT PAGES (0-63) OF VIRTUAL MACHINE NUMBER 1

INPAG	OUTPAG	INSTR	*****			
013000	008000	002674	++++++	+ + + + 0		+
014000	001000	002674	+ + + + 0000	0 + + + + 0		0
015000	009000	002674	+ + + + 00	0 + + + + 0		0
016000	008000	002674	+ + + + 0	0 + + + + 0		0
01A000	00E000	006E4C	+ + + + +	+ + + + +	0	+
01B000	002000	013C4C	+ 0 + + +	0 + + + + 0	+ 0	0
002000	010000	002760	+ 0 + + +	+ + + + 0	+	0
010000	03F000	0005AC	+ + + + +	0 + + + + 0	+	
01C000	006000	013410	+ + + + +	+ + + + + 0	+	
01D000	005000	013410	+ + + + +	+ + + + + 0	+	
01E000	011000	013410	+ + + + +	+ + + + 0	+	
01F000	007000	013410	+ + + +	+ + + + 0	+ + + + +	
020000	003000	013410	+ + +	+ + + + 0	+ + + + +	
021000	004000	013410	+ +	+ + + + 0	+ + + + +	
022000	012000	013410	+ +	+ + + + 0	+ + + + +	
023000	016000	013410	+ +	+ + + +	+ + + + + + 0	
024000	002000	013410	+ +	0 + 00	+ 00000000 + 0	
025000	01D000	013410	+ +	0 + 00	+ 00 0000 + 0	
026000	014000	013410	+ +	0 + 0	+ 00 0000 + 0	
027000	01E000	013410	+ +	0 + 0	+ 00 0000 + 0	
028000	015000	013410	+ +	0 +	+ 00 0000 + 0	
029000	01C000	013410	+ +	0 +	+ 000 + + + + + 0	
02A000	01F000	013410	+ +	0 +	+ 000 + + + + + 0	
012000	020000	01345A	+ +	0 0 +	+ 00 + + + + + +	
007000	021000	007B3C	+ +	0 + +	+ 0 + + + + + +	
005000	01B000	005178	+ +	0 + +	+ 0 + + + + + +	
011000	022000	00549A	+ +	00 + +	+ + + + + + +	
002000	010000	0020A8	+ +	0 + +	+ + + + + + +	
003000	013000	0020B6	+ +	00	+ 00000000	
010000	01A000	0005AC	+ +	000	+ 00000000	
01A000	023000	002682	+ +	+ + 0	+ 00000000	
013000	025000	013798	+ +	+ + +	+ 0 00000	
006000	026000	006C40	+ +	+ + +	+ 0 000 +	
004000	027000	003546	+ +	+ + + +	+ 0 00 +	
00E000	028000	00E688	+ +	0 + + + +	+ 0 0 +	
00F000	029000	000752	+ +	+ + + + +	+ 0 +	
014000	024000	002674	+ +	+ + + + + 0	+ +	
015000	02A000	002674	+ +	+ + 0 + + 0	+ 0	
016000	012000	002674	+ +	+ + 0 + + 0	+ 0	
03F000	007000	002674	+ +	+ + 0 + +	+ 0	
012000	002000	012F38	+ +	+ 0 + +	+ 0	
01B000	003000	015634	+ +	0 + +	+ 0 + + + + + +	
007000	010000	007B3C	+ +	0 + +	+ 0 + + + + + +	
002000	004000	0020A8	+ +	0 + + + +	+ 0 + + + + + +	
003000	00E000	0020B6	+ +	+ + + + +	+ 0 + + + + + +	
010000	01B000	002682	+ +	00000000	+ 0	
02A000	013000	0168F8	+ +	0 + + 00 +	+ 0	
013000	006000	012FFE	+ +	0 + + 000 +	+ +	
01B000	00F000	015634	+ +	+ + + + +	+ 0	
02B000	03F000	0168F8	+ +	+ + + + +	+ +	
004000	01A000	0043A0	+ +	+ + 0 + 0	+ 000000	
01A000	015000	012CE6	+ +	+ + + + +	+ 00	
015000	01B000	01573A	+ +	+ + + + +	+ +	
01B000	004000	015634	+ +	+ + + + +	+ 0	
004000	010000	0043A0	+ +	+ + 0 + 0	+ 000000	
010000	013000	002F2C	+ +	0 + 0 000	+ 00	
013000	014000	013B2C	+ +	+ + + + +	+ 0	
014000	02B000	014DA4	+ +	+ + + + +	+ 0	
02B000	002000	0168F8	+ +	+ + + + +	+ +	
002000	003000	0020A8	+ +	+ + + + +	+ +	
003000	02A000	0020B6	+ +	+ + + + +	+ +	
02A000	01A000	002AAA	+ +	+ + + + +	+ +	
01A000	016000	012CE6	+ +	+ + + + +	+ +	
016000	000000	016874	+ +	+ + + + +	+ +	
000000	011000	005178	+ +	0 + + 0 + +	+ +	
011000	012000	00549A	+ +	0 + + 0 + +	+ +	
012000	002000	007AF4	+ +	0 + 0 + 0 + +	+ +	
002000	014000	0020A8	+ +	0 + + + +	+ +	
014000	004000	014DA4	+ +	+ + + + +	+ +	
004000	02B000	0043A0	+ +	+ + + + +	+ +	
01C000	02A000	016210	+ +	+ + + 0 + +	+ + 0	
01D000	013000	016210	+ +	+ + + 0 + +	+ + + 0	
01E000	01B000	016210	+ +	+ + + 0 + +	+ + + +	
01F000	003000	016210	+ +	+ + + 0 + +	+ + + + +	
020000	010000	016210	+ +	+ + 0 + +	+ + + + +	
021000	007000	016210	+ +	+ + 0 + +	+ + + + +	

022000	015000	016210	+ + + +	++ 0 +	+ + + + + + +	
023000	005000	016210	+ + +	++ 0 +	+ + + + + + +	
024000	011000	016210	+ + +	+ 0 +	+ + + + + + +	
025000	014000	016210	+ + +	+ +	+ + + + + + + + 0	
026000	002000	016210	+ 0	0 +	+ 000000000+0	
027000	012000	016210	+ 0	0 +	+ 00 000000+0	
028000	017000	016210	+ 0	0 +	+ 00 00000+0	
029000	020000	016210	+ 0	0 +	+ 00 0000+0	
02A000	021000	016210	+ 0	0 +	+ 00 000+0	
015000	004000	015968	+	0 0+	+ 00 000+0	
007000	022000	007A80	+ 0	+ ++	+ 00 00+0	
005000	01C000	005178	+ 0 +	+ ++	+ 0 00+0	
011000	023000	00549A	+ + +	0+ ++	+ 0 0+0	
003000	01D000	0033D8	+ 0 + +	+ + ++	+ 0 +0	
01B000	024000	015990	+ + + +	+ + ++	+ 0 +0	
024000	016000	015A9E	+ 0 0 0	00 +	+ + 0+0	
023000	012000	015A9E	+ 0 0 0	0 +	+ + 0+0	
022000	000000	015A9E	0 0 0	0 +	+ + +0	
021000	007000	015A9E	0 0	0 +	+ + 0+0	
020000	005000	015A9E	0	0 +	+ + 0+0	
01F000	011000	015A9E	0	+ +	+ + 0+0	
01E000	003000	015A9E		+ +	+ + 0+0	
01D000	01B000	015A9E		+ +	+ 0+0000000000	
01C000	027000	015A9E		+ +	+ 0+0000000 000	
01B000	028000	015A9E		+ +	+ 0+0+000000 00	
028000	025000	015A9E		+ +	+ + + + + 00000 0 0+	
027000	024000	015A9E		+ +	+ + + + + 0000 00+	
025000	023000	015A9E		+ +	+ + + + + 000 0+	
024000	022000	015A9E		+ +	+ + + + + 00 0+	
023000	029000	015A9E		+ +	+ + + + + 00 + + + + +	
022000	02A000	015A9E		+ +	+ + + + + 00 + + + + +	
012000	026000	012BC8		0 +	+ + + + + + 0000 00	
007000	015000	007748		+ +	+ + + + + + 0000 00	
005000	01A000	005178	+ +	+ +	+ + + + + + 0000 00	
000000	028000	005178	0 + +	+ +	+ + + + + + 0000 0	
011000	023000	00549A	+ + +	0+	+ + + + + + 00 0	
002000	027000	0020A8	+ 0 + +	+ +	+ + + + + + 00	
003000	022000	0020B6	+ + 0 + +	+ +	+ + + + + + 00	
00F000	025000	0005AC	+ + + +	0 + +	+ + + + + + 0	
010B80	024000	000CF2	+ + + +	+ 0 + +	+ + + + + +	
02A000	01C000	002674	+ + + +	+ + + +	0 00000	0
01A000	01B000	012CE6	+ + + +	+ + + +	0 00000	+
015000	021000	015960	+ + + +	+ + + + 0	+ 0000	+
016000	020000	016362	+ + + +	+ + + + + 0	+ 000	+
013000	01F000	013AC8	+ + + +	+ + + + 0 +	+ 00	+
004000	01E000	0043A0	+ + 0 + +	+ + + + +	+ 0	+
02B000	01D000	002AAA	+ + + +	+ + + + +	+ +	+ 0
006000	015000	006C40	+ + + + 0	+ + + + 0	+ +	+
014000	027000	007674	+ + + + 0	+ + + + 0 0	0	0