MEMORY EXTENSION TECHNIQUES FOR MINI-COMPUTERS

by

Mary Poppendieck
and
Edward J. Desautels

Computer Sciences Technical Report #290

December, 1976

MEMORY EXTENSION TECHNIQUES FOR MINI-COMPUTERS

by

Edward J. Desautels
and
Mary Poppendieck

Computer Sciences Department
University of Wisconsin
December, 1976

INTRODUCTION

There are many computer users these days who have bumped up against the end of their memories. Years ago, IBM 1800, 1130, and 7094 users discovered that 32K words of addressable memory simply was not enough. Today, mini-computer users are discovering the same thing. And so are mini-computer manufacturers. Some of the newest mini's can address 128K, 256K, 512K or even 1024K (1M) bytes of main memory. The claim is that 16M byte addressing is right around the corner.

It makes sense to add more and more memory to mini-computers, because memory prices are dropping rapidly. In assessing the future cost of memory, one must distinguish between the memory device components (e.g. memory chips) whose costs are dropping more rapidly, and the memory systems these go into (e.g. memory chips and supporting logic, circuit boards, etc.). The component cost for 4096 bit chips was in the $3.50 to $4.00 range in late 1976, or approximately 0.1¢/bit, whereas the memory system cost to the system assembler was approximately 2.6 times greater (.26 ¢/bit). This factor is expected to decrease by 1980 to approximately 2.0 for smaller memory system (512K bytes for $1300) and to 2.5 for larger memory systems (4 MW, 64 bits/word, with error correction circuitry, for $100,000). The cost to the end-user is likely to involve another factor of 2 or 3. Figure 1. depicts the cost/bit at the component and system levels over time. As of January, 1976, it was possible for an end-user to obtain a 16K 16-bit word

memory board for less than $1300 in unit quantity, or 0.5¢ per bit. When this is compared to the salary of a programmer trying to squeeze code into a small memory, extended memory becomes very attractive now, and it will be even more so in the future.

However, the mini-computer user contemplating memory expansion must understand the various memory expansion techniques which are available in order to evaluate their effectiveness. One computer's 1M bytes is not the same as another computer's 1M bytes. There are dramatically different ways of addressing large memories.

It is important to establish the fact that there are two kinds of limits to memory space, logical and physical. Logically, a program is limited to a certain address space by the instruction set of the machine. Any linked set of code is usually confined to a certain logical space, depending upon the address size supported by the instruction set. Most mini-computers allow a maximum of 16 bits for an address, which gives 64K addresses. However, in some computers this address is a byte address, and if the computer's word size is 16 bits, only 32K words can be addressed. If one is dealing with applications in which floating point values are required, then the 32K of 16 bit words becomes a maximum of 16K floating point words of 32 bits. Many other mini-computers reserve one of the 16 bits as a flag, perhaps for indirect addressing, which again limits the logical address space to 32K words. If a single program with its data cannot fit within the logical address space, more memory will not help unless the program is divided into segments or overlays.

One method of extending mini-computer memory is to extend the logical address space by changing the classical instruction format and allowing more than 16 bits for an address. Currently the Interdata 7/32 and 8/32 instruction sets allow a 24 bit address and the Scientific Engineering Laboratories SEL32 instructions allow 19 bits for an address. These computers are basically 32 bit computers at the top of the "mini" classification.

The physical limits of the memory space are determined by the number of wires carrying an address to the memory, as well as the granule size which is selected by the address. For example, in a PDP-11/45, eighteen wires carry the address to memory, allowing 256K addresses. But since each address is a byte address and there are two bytes per word, only 128K words may be physically addressed, with the top 4K reserved for access to peripheral devices. The logical address space of a PDP-11 is 32K words, with 4K of that dedicated to the bus hardware.

If memory extension beyond the logical address space is implemented, then the computer can no longer simply use the logical address to select a physical address; some transformation or translation must be applied to the logical address to convert it to a physical address. This translation consists of combining the logical address, or some portion of that address, to a base address which is obtained in some specific way. The method of selecting the base address and combining it with the logical address can be used to classify physical memory extension.

techniques. There are three basic methods of implementing the translation of a logical address into a physical address: memory mapping bank switching, and the use of base registers.

MEMORY MAPPING

Memory mapping involves dividing the logical address into two parts-a descriptor, which is a set of the most significant bits, and a displacement, which is the remaining set of bits of the address. A function is applied to the descriptor, and the result of this function is a physical address which may be called a page address. The displacement is added to the page address to give a physical address. Generally, pages in memory do not overlap, and the number of bits in the displacement determine the granularity, or page size, of the memory. Typical page sizes vary between 256 and 1024 words. In most cases the significant part of the page address is effectively concatenated with the displacement. (See Figure 2a). In some implementations (e.g. PDP-11/45 and TI990/10) the page size is smaller than the displacement field, so that the displacement field must be subdivided into a second descriptor and the second displacement. The second descriptor is added to the beginning page address to obtain the true page address, which is added to (or concatenated with) the second displacement to obtain the physcial address. This technique gives the PDP-11/45 and 11/70 a page size between 32 words and 4K words. But, their limited number of mapping registers makes it necessary to use 4K pages in order to gain sufficient address space. (See Figure 2b).

ASSOCIATIVE MAPPING

Memory mapping was first used in the early sixties. The Atlas computer developed at Manchester University used a memory mapping technique which was eventually adapted for use with associative memories and became the prototype of several memory mapping systems. Memory was divided into 32 "pages" of 512 words, and there were 32 Page Address Registers, one for each physical page. The page address registers were directories to the program. As a segment of the program was brought in from the drum being used as the pageing device, the program segment number was filed in the page address register corresponding to the physical page where the program segment was located. If a reference was then made to that program segment, the computer proceeded to search through the address registers, looking for the program segment number. If a match was found then the physical page number of the memory reference was simply the register number of the match. If no match was found, the missing program segment was brought in off a drum. (See Figure 3a).

A sequential search of several registers for every memory reference can be slow, but content addressable or associative memories remove this drawback. A key and a value are stored in one slot of a small associative memory (on the order of eight registers are usually used). The key is the descriptor from the logical address and the value is the physical page. When a descriptor is sent to the associative memory, all keys are simultaneously searched. If a match is found, the value,

a physical page number, is the result of the compare. If no match is found, the associative registers are updated from a table in memory. (See Figure 3b). This type of map was first used in the IBM 360/67 (1966) and in Scientific Control Corporation 4700 mini-computer (1968). Today it is found in the Prime Computer Models 300 and 400.

INDEX MAPPING

Another type of memory map was used on the Berkeley Timesharing system SDS 930, about 1965. In this system there were 8 six-bit registers which could be selected by the top 3 bits of a logical address, used as an index. The contents of the selected register contained a six bit physical memory page number. (See figure 4a). A memory map of this type is available from Fabri-tek for the PDP-11/05, 11/10, 11/15, and 11/20. Note that the top three bits of each address will always select a register- there is no possibility of a no-match. Therefore an entire program must be in memory before execution can begin, unless there is some way of telling the system whether or not each register contains a valid address.

If protection bits of some sort are added to the registers the requirement that an entire program must be in memory before execution is removed from the index-type mapping scheme. Protection bits are usually used to restrict users from reading, writing into, or executing certain pages of physical memory. However, since reference to a completely protected page will cause a trap of some sort, the protection bits can be used to indicate that a page is not in memory and must be

fetched from mass storage. An extra protection bit can be used during a program's execution to determine whether or not a page has been changed and thus whether or not the page must be saved before it is overlaid. (See Figure 4b).

In contrast to associative mapping registers, registers which are addressed by the top bits of a logical address used as an index are relatively inexpensive. With an associative map, the various maps are usually stored in main memory and when a memory reference creates a "no-find" condition, the associative registers are updated from main-memory. With an index-type map, the map or maps are stored in the mapping hardware registers. Several sets of registers may be available. The HP21MX has two maps - a system and a user map. The PDP 11/45 and 11/70 have three maps - kernel, supervisor, and user. The Modcomp IV has four maps - one for system, two for large users, and one which may be subdivided among four small users. The Data General Eclipse has three user maps and the supervisor is not mapped. The IBM Series 1 has eight maps. The Varian 75 provides sixteen maps in 1024 registers. The Harris Slash 4 map provides 1024 registers for any number of users, since the beginning of the active map is selected by a 10 bit base register.

More hardware maps allow faster switching between users. Some computers, like the PDP-11/45, 11/70 and the Modcomp IV, divide the maps up into an instruction and a data map. If the data and instructions can be completely separated, the address space available to any one

program can be doubled, since the same addresses could be used for instructions and data, and then the map can separate the two. (See Table 1). The mappings on the PDP-11/34, 40 and 55 do not distinguish between instructions and data.

EVALUATION

The limitations on logical address space in any mapped system is the most serious drawback of any mapping scheme. No matter how large the physical addresses can become, the logical address space remains limited by the computer's instruction set. The only way around a limited logical addresss space is program segmentation and overlays; more physical memory, if it is mapped, won't solve the problem of a program which has outgrown its logical address space.

Another very important drawback of memory mapping is I/O handling. In general, I/O does not occur through the memory map. Even if most I/O could go through the map, DMA (direct memory access) channels from a disk, say, must access large sections of data quickly, with no time for a "no match" condition to occur. Moreover, if one user asks for I/O, control may very well switch to another user, so that the user map would not correspond to the I/O map. In most systems, therefore, I/O is handled using physical memory addresses. If block I/O is going to physical addresses, but the I/O buffer is not in consecutive locations in physical memory (as can easily happen with mapping) it is probably necessary for a supervisor program to do all I/O through its own buffers and copy the buffers from or to the user buffer. Whereas this might be

reasonable in a timesharing system, it can be a severe time restriction in a dedicated system, especially one with the large data buffers that are common in real time monitor systems.

Some computers have some of the I/O going through a separate map. The Eclipse has one data channel map, the HP21MX has two such maps. The PDP-11/70 has a Unibus map. To be useful for large data buffers, these maps must apply to DMA access. The high speed I/O channel on the PDP-11/70 is not mapped, thus high speed transfers of data across block boundaries cannot be easily handled. The IBM Series 1 allows each DMA device to transfer through any of the maps. When I/O is performed using physical addresses, the supervisor may not, itself, be mapped, or it may be mapped into consecutive pages in memory. This allows the supervisor to determine the physical addresses of its own buffers, and assures that these buffers are in contiguous blocks of memory. A supervisor which is not mapped can save mapping overhead time and simplify the "startup " problem, namely that map tables must be initialized by the supervisor before any mapping of the supervisor can begin.

A third problem with memory mapping is a slight decrease in speed. The PDP-11/45 map adds 90 ns to each memory reference; the Varian 75 map adds 30 to 40 ns. Manufacturers claim that the HP21MX and PDP-11/70 maps do not add to the memory reference speed. This may be true for today's memories, but probably not for the higher speed memories of tomorrow. Besides the actual hardware construction of each address, the mapping unit must be set up with each user's map. The maps must be switched as

new users are enabled. Associative registers require an extra memory reference with each "no match" condition. Unmapped I/O which must cross page boundaries can dramatically decrease "high speed" I/O time. And finally, if a segment of a program must be fetched from mass storage, a large swapping time will be added.

Memory mapping usually provides a supervisor and user mode, as well as page level memory protection and traps on certain instructions attempted while in user mode (such as halts and I/O instructions). It allows for multiple users of various sizes. It also requires a layer of software and a layer of hardware to be added to the existing system. Both the software and the hardware must be of a rather sophisticated design and they carry a relatively high initial cost. Moreover, this software and hardware may degrade the running time of existing programs to levels which are unacceptable, especially in a real time environment.

BASE REGISTERS

Using base registers involves adding the logical address to an address found in a base register. (Figure 5) Again, there may be two registers, one for instructions and one for data references. With base registers the granularity of memory is not usually a consideration, since the base register schemes tend to limit the logical address space by limiting the displacement which is added to the base register. In this case, the logical address space is expanded by changing the base registers. The IBM 360's and 370's use base register changing to

determine all addresses, but there do not seem to be any mini-computers which use base registers to determine the addresses for every memory reference. This is probably because of the sixteen bit word orientation of most mini-computers. Base registers must be as wide as the physical addresses the memory system require, perhaps twenty or twenty-four bits wide, but sixteen bit machines do not usually support registers wider than sixteen bits.

Although base registers are not used to extend mini-computer addressing capabilities, they are found in some mini-computers. A special case of base register usage occurs when the program counter doubles as a base register.

Many mini-computers add a signed displacement field to the program counter to determine some operand addresses. The Data General Nova has an 8 bit displacement; the Interdata 7/32 allows 15 bits. By using the program counter as a base register, programs are not restricted to particular segments of physical memory. Instead, the direct address range of an instruction is the memory area surrounding the instruction as limited by the displacement size. This technique allows direct addressing beyond either end of a program if the displacement permits, but it also requires careful program organization if the program is larger than the displacement range.

BANK SWITCHING

Bank switching involves concatenating a bank indicator, which is separate from the logical address, with the entire logical address to obtain a physical address. In some machines, two separate bank indicators are used for instruction and data memory references. In any case, since the entire logical memory address determines the displacement within a page, the memory granularity is quite large. (See Figure 6).

Bank switching had its roots in the early 1960's, on a few IBM 7094's and DEC PDP-5. About 1962, MIT added an extra 32K to the 7094II used for their Compatible Time Sharing System (CTSS). A single bit register was used to select which 32K bank memory was in use. In addition to expanding the memory, this system prevented programs in one half of memory from interfering with programs in the other half. About the same time, the PDP-5, predecessor of the PDP-8, was able to address 32K of memory with only 12 bits for the address. This was done with a set of two three-bit registers which could be read and written with special instructions. These registers selected one of eight banks of 4096 words. One register selected the instruction bank and one selected the operand (data) bank. This technique is still used by the PDP-8 today.

METHODS OF BANK SWITCHING

Computer Automation's Naked Mini and Megabyter LSI-2 Computers can address 512K and 1M bytes respectively, using a bank switching option. The logical address space is only 32K. General Automation's SPC-16 and the Microdata 1600 also have a bank switching option.

There are many variations of bank switching available. The SPC-16 divides its 32K address space into eight 4K banks. The lower four banks are always the same, but the upper four banks may be switched to any of eight additional 4K banks which are selected by an I/O command. This provides 64K of physical memory, with 32K selected at any one time. In one system developed by the Nuclear Physics group at the University of Wisconsin, a Honeywell DDP-124 has a flip flop on the end of an I/O channel acting as a bank switch. The lower 8K of each bank is physically the same, and in this 8K resides a monitor which communicates with the 24K left in either bank and does all the bank switching. With this system, a total of 16K of address space is lost to the monitor.

Harris Slash 5 computers can directly address 32K 24 bit words with most instructions, but they support 64K of memory. In this case the top bit of the address comes from the program counter. A special jump instruction has a 16 bit address field for access to the other bank, or else indirect addressing expands the address space to 32K.

The top bits of the address can also be obtained from the status register. For instance if there is both a system and a user mode which is indicated in the status register, the system/user mode bit can be used to determine the top bit of the address. In the modified PDP-11's used by the Carnegie-Mellon multi-mini processor, a two bit bank indicator is contained in the program status word. This gives the advantage of saving the bank indicator automatically whenever the status word is saved.

EVALUATION

Three distinct disadvantages of bank switching should be mentioned. First of all, communication between banks may be very difficult, if not possible. This may be an advantage, since memory protection increases as cross-talk becomes more difficult. However, if a supervisor is to run programs in all banks, some method of communication is usually needed. Usually communication between banks must be done via registers or disk. In the SPC 16 and DDP-124 systems described above, communication is simplified because the lower quarter or half of memory remains unchanged.

The second serious drawback of bank switching is the granule size. The banks are typically large and the logical space of any program must usually be confined to one bank. This is a rather rigid restriction which would generally result in a lot of wasted memory space. Moreover, memory protection within each bank may not be available. However, in some systems, particularly dedicated ones, there are often only two users of

approximately equal sizes, and two banks. Or it may be feasible to separate the supervisor and the user into separate banks, especially if the supervisor is large or can be assigned enough tasks to make it into a second user.

The third serious drawback of bank switching is that banks must be switched. First of all, what is going to happen to the program counter when the banks are switched? This is not a problem when the program counter is also the bank indicator, but is definitely a problem otherwise. If there isn't some sort of memory common to both banks in which the bank switching can occur, then there must be some way of changing the program counter at the same time as switching the banks. Secondly, how is the switching mechanism governed? Switching from one bank to another usually requires positive action on the part of the user or supervisor. Protecting the program status word from the user (if it contains the bank indicator) or not allowing the user access to the bank register is probably necessary. Some method of saving the current bank in case of an interrupt and restoring it after the interrupt is needed. And of course, some means of deciding when to switch from one bank to another must be available.

BANK SWITCHING FOR SPEED

In a dedicated system with only a few users, or a single user and a supervisor, bank switching may be more economical and simpler technique for extending memory. Table 2 gives an overall picture or the tradeoffs involved in using bank switching vs. memory mapping for such a system.

One possibility, implemented by the High Energy group of the Physics department at the University of Wisconsin on a Scientific Control Corporation 4700 with a 32K logical address space, is a hybrid scheme using some features of memory mapping in a bank switching environemnt. In this implementation,real time considerations ruled out the use of a fully mapped system, although mapping hardware was available. It was estimated that the associative map would add about 10% to execution time. Furthermore, the transfer of large data buffers (usually 2K) through DMA units was necessary, and these could not be copied by the supervisor without an intolerable degradation in speed. Thus it would be necessary to map all user pages consecutively, if mapping were to be used. Instead, the system was divided into two 32K banks, with one bank dedicated to the supervisor and one to the user. A large and independent data handling routine became part of the supervisor, giving two approximately equal size programs. The mode bit in the status register (system or user) determines which bank is in use. Bank switching is accomplished by a system call instruction (switch to one of 64 system routines) and system return (return to caller) instruction normally used by the mapping hardware. These instructions simultaneously switch the system/user bit in the status register and save or restore all registers, including the status register. Communication between supervisor and user can take place because the supervisor can place itself in an "indirect user" mode. In this mode, the last memory reference of an indirect address sequence is done in user mode, while all other memory

references are in system mode. This allows the supervisor to read or write in the user bank under special circumstances. Thus in this implementation, a few features of a memory map system are being used to overcome drawbacks of bank switching, resulting in a very efficient bank switching system.

SUMMARY

Probably the most straight forward approach to extending mini-computer memory is by extending the number of address bits which can be placed in an instruction. However, this technique cannot be used by manufacturers with long established instruction sets, nor is it being used by every manufacturer coming out with a new machine. Instead of departing from the traditional instruction set design mini-computer manufacturers are using either bank switching or memory mapping to extend the physical size of their computers' memory.

Bank switching is the more primitive approach, involving an action by the program in order to switch banks. It is likely to be faster and less complex than memory mapping, and may be the method of choice for dedicated, real time use.

Memory mapping is more complex than bank switching, and provides the kind of context switching and memory protection needed for time sharing. In fact, even machines with a large direct address space, such as the Interdata 8/32, provide memory mapping capabilities. Memory mapping implementations vary greatly in sophistication, and potential users may wish to inspect the number of map files provided and insure that

DMA paths go through a map.

Above all, a mini-computer buyer should not assume that programs of any size may be run in a mini-computer, simply by expanding the memory. The data space requirements of programs tend to outgrow the addressing capabilities of mini-computers fairly quickly in some applications, so that it is usually necessary to change the addressing capability of the instruction set in order to gain more space for an individual program. This kind of expansion cannot be accomplished with either memory mapping or bank switching. It takes a new instruction set format, and therefore, a new computer.

BIBLIOGRAPHY

1. Bell, C.G., and A. Newell, Computing Structures: Readings and Examples, McGraw-Hill, New York, 1971.

2. Doran, R. W., "Virtual Memory," Computer 9, 10(Oct., 1976) pp.27-37.

3. Freeman, P., Software Systems Principles-A Survey, Science Research Associates, Chicago, 1975.

4. Hobbs, L.C. and R.A. McLaughlin, "Minicomputer Survey," Datamation 20, 7(July, 1974), pp. 50-61.

5. Hodges, D. A., A Review and Projection of Semiconductor Components for Digital Storage," Proc. IEEE, Vol. 63, No. 8, Aug. 1975, pp.1136-1147.

6. Hollingworth, D., Minicomputers: A Review of Current Technology, Systems, and Applications, Rand Corporation, Santa Monica, California, 1973.

7. Kilburn, T. et al, "One Level Storage System," IRE Trans., EC-11, vol. 2 (April, 1962), pp. 223-235.

8. Koppel, R.J. and I. Maltz, "Predicting the Real Costs of Semiconductor-Memory Systems," Electronics, Nov. 25, 1976, pp. 117-122.

9. Lampson, B.W. et al, "A User Machine in a Time-sharing System," Proc. IEEE, 54, vol. 12 (December, 1966), pp. 1766-1774.

10. Michaels, John, "The Mega-Mini Succeeds the Model T," Datamation 20, 2 (February, 1974), pp. 71-74.

11. Reagan, Fonnie H., Editor, Datapro Reports on Minicomputers, Datapro Research Corporation, Delran, N. J., 1975.

12. Watson, R.W., Timesharing System Design Concepts, McGraw-Hill, New York, 1970.

Model	Basic Unit	Directly Addressable Address Space	Indirectly Addressable Address Space	Maximum Physical Memory	Type of Address Translation	Page Size	Number of CPU maps /System	Number of CPU maps /User	Number of I/O Maps
Naked Mini	8 or 16 bits	32Kw	32Kw	256Kw	Bank Switch		--	--	--
Megabyter	8 or 16 bits	32Kw	32Kw	512KW	Bank Switch		--	--	--
Nova 3/12	8 or 16 bits	1Kw	32Kw	128Kw	Index Map	1Kw	0	2	2
Eclipse	8, 16, or 32 bits	1Kw	32Kw	128Kw	Index Map	1Kw	0	2	1
ols D-616	16 bits	1Kw	32Kw	1Mw	Index Map		--	--	--
PDP-8	12 bits	256Kw	4KW	32Kw	Bank Switch	4Kw	--	--	--
PDP-11/45	8 or 16 bits	32Kw	32Kw	124Kw	Index Map	32w - 4Kw	2*	2*	0
PDP-11/70	8 or 16 bits	32Kw	32Kw	2Mw	Index Map	32w - 4Kw	2*	2*	1+
Solution 16/440	8 or 16 bits	64Kw	64Kw	1Mw	Index Map		1	1	2
Series 200	24 bits	32Kw	192Kw	256Kw	Index Map	1Kw		Variable	0
21MX	8, 16, or 32 bits	2Kw	32Kw	1Mw	Index Map	1Kw	1	1	2
Series 1	8 bits	64Kb	64Kb	16Mb	Index Map	2Kb		-see note** -	
8/32	8, 16, or 32 bits	1Mb	--	1Mb	not needed	--	--	--	--
ms Modcomp IV	16 or 32 bits	64Kw	64Kw	256Kw	Index map	256w	1	3-8	0
Prime 400	16 bits	64Kw	64Kw	8Mb	Asso. Map				
SEL 32	32 bits	1Mb	1Mb	1Mb	not needed	--	--	--	--
990/10	8 bits	64Kb	64Kb	1Mw	Index Map	16w - 32Kw	1	1	--
V75	16 bits	32Kw	32Kw	256Kw	Index Map	512w	1	15	0

Table 1: Memory Expansion in Representative Mini-Computers. Information from Datapro Reports on Minicomputers and Vendor-Supplied Literature.

ne data map
. Each priority level
instruction and two data maps
-type device may be assigned

K = 1024
M = 1024K
w = word
b = byte

Memory Mapping				Bank Switching			
Software	IO	Hardware		Software	IO	Hardware	
2	3	2		3	3	3	Low Implementation Cost
2	1	2-3		4	4	4	High Speed
3	2	3		3-4	4	4	Minimum Memory Used for Implementation
3-4	2	3-4		1	3	1	Minimum Memory lost to Fragmentation
4	4	3-4		1	1	1	Large Number of Users in Memory
3	2	4		2	4	1	Transparency to User
3	2	3		1-3	3	1-3	Communication between Supervisor & User

4 = Very Good
3 =
2 =
1 = Very Poor

Table 2: - Bank switching vs. memory mapping for a dedicated mini-computer. Each row should be weighted for the importance to the application. Grids with variable scores are highly dependent on the implementation.

Mary Poppendieck is currently a Project Engineer with General Motors Corporation Transportation Systems Division, in Warren, Michigan. Work on this paper was done while she was a Systems Analyst in the Physics Department at the University of Wisconsin, Madison, Wisconsin, working on real time equipment control and data aquisition.

She received a Masters Degree in Mathematics from the University of Maryland in 1971, and a B.S. from Marquette University in 1967. Ms. Poppendieck has been an active member of ACM since 1972 and has been a member of IEEE Computer Society since 1976.

***

Edward J. Desautels is an associate professor in the Computer Sciences Department of the University of Wisconsin-Madison. He is the director of the department's computer laboratory, and the department's associate chairman. His research and teaching have involved programming systems, languages and applications.

He received a B.Sc. from the University of Manitoba, an M.Sc. from the University of Ottawa, and was awarded a Phd.D. in Computer Science from Purdue University in 1969.

Dr. Desautels is a member of the IEEE and the ACM. He recently chaired an ACM panel on computing facilities for computer science education, and is chairman of the ACM Special Interest Group on Computing and the Physically Handicapped.