LBL-13494

Lawrence Berkeley Laboratory

UNIVERSITY OF CALIFORNIA

Engineering & Technical Services Division

Submitted for the Texas Instruments Members' Information
Exchange National Symposium, Las Vegas, NV, March 7-10, 1982

A PROPOSED MIDAS II PROCESSING ARRAY

John Meng

March 1982

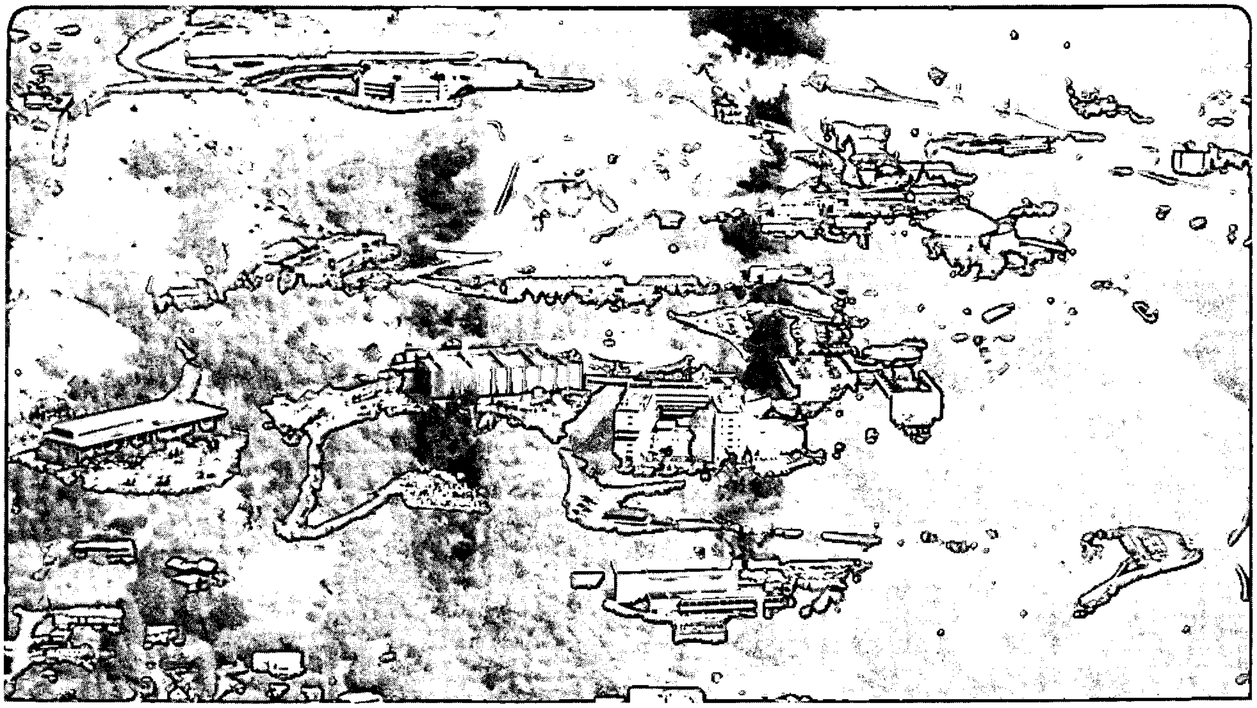

Prepared for the U.S. Department of Energy under Contract DE-AC03-76SF00098

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California.

LBL-13494

A PROPOSED MIDAS II PROCESSING ARRAY

John Meng

Lawrence Berkeley Laboratory
University of California
Berkeley, California USA 94720

Abstract

MIDAS (Modular Interactive Data Analysis System) is a ganged processor scheme used to interactively process large data bases occurring as a finite sequence of similar events. The existing device uses a system of eight ganged minicomputer central processor boards servicing a rotating group of 16 memory blocks. A proposal for MIDAS II, the successor to MIDAS, is to use a much larger number of ganged processors, one per memory block, avoiding the necessity of switching memories from processor to processor. To be economic, MIDAS II must use a small, relatively fast and inexpensive microprocessor, such as the TMS 9995. This paper analyzes the use of the TMS 9995 applied to the MIDAS II processing array, emphasizing computational, architectural and physical characteristics which make the use of the TMS 9995 attractive for this application.

Introduction

MIDAS (Modular Interactive Data Analysis System) is a pyramidal processing scheme (Fig. 1) whereby the peak of the pyramid is the greatest hardware and software "intelligence" in the system, passing instructions downward through progressively increasing multiplicities of progressively less "intelligent" processors. Interactive user stations connect to the top of the pyramid. A massive data base connects to the bottom. The functional objective of the system is the multiple scanning and processing of the entire data base at rates high enough to make the user an interactive part of the data analysis.

Raw data from the data base passes through an array of processors, all possibly executing the same code. Data flows from the data base into one of 16 array memories through a pipeline input processor. The memory is then connected to one Central Processing Unit (CPU). When a data processor, consisting of an array of minicomputer CPU's, finishes the processing of the memory load of data, the memory is switched onto an output processor which sends the processed data to a preselected destination. Operation of the prototype MIDAS I containing only three CPU's has shown it to be practical to process a 9-track 800bpi, 2400 ft. reel of tape in less than a minute. MIDAS has been described in great detail in available publications, and the prototype of one processing segment is operating.

We are presently considering the next level of implementation, MIDAS II, in which we will use up-dated technology to make it even more economic and effective than the prototype. The modular character of MIDAS allows the simultaneous inclusion of radically different processing modules within the pyramid, so our next generation may contain processors from two years ago working side-by-side with those designed using today's or tomorrow's technology. Within the pyramid (Fig. 1) we will be discussing the "supermicro" level, containing eight central processors from a high performance minicomputer. Since data flows into and out of this level by physically gating processors onto memories already containing data, no input/output in the conventional sense is required. The eight data processing units share sixteen memories. As processing modules become faster and less expensive, it may be practical to assign each memory its own dedicated processor. Towards this end, we have considered using the TMS 9995. MIDAS lends itself to many applications. Although there are some for which the use of a micro in the data processing array would be inappropriate, there are some for which the dedicated micro-per-memory configuration has substantial benefits. Once again, the MIDAS configuration allows for the inclusion of many different types of processors simultaneously working side-by-side within the pyramid.

Sizing Up the Competition

TAble 1 compares the TMS 9995 running Power Basic with the TMS 9900 running Power Basic. Such a comparison is of interest because, although the 8-bit data bus on the TMS 9995 will tend to slow it down by requiring additional memory accesses, it executes instructions in less time (see Table 2). The TMS 9995 runs 10-20% faster than the TMS 9900 in this test. High speed memories in the TMS 9995 module would have allowed the elimination of the memory wait during memory accesses, further enhancing the differences between the two. The Power Basic test loops are similar, except the first loop is computation-intensive, the second contains only one simple computation, and the third does a 16-bit CRU input with no computation. Time measurements were done with a stopwatch while listening for the printer to start. All three loops executed in significantly less time on the TMS 9995, and the speedup is 10-20%.

Table 2 is a book-derived comparison among three devices. The prototype MIDAS I is operating with minicomputer CPU boards to do the processing. Many applications may allow the use of a microcomputer, such as the TMS 9995, to be substituted for the minicomputer CPU cards. Thus, we have compared several characteristics of the currently used minicomputer CPU with those of a (computer-wise) comparable TMS 9995 system. Times are included for a TMS 9900 (990/101 board) to complete the picture in the comparison of the TMS 9995 with the TMS 9900. The numbers are derived from manufacturers' published descriptions. A single wait-state is assumed for TMS 9995 and TMS 9900 memory accesses. Comparisons of this nature tend not to be accurate predictions of performance in specific applications. For example, the TMS 9900 executes, on the average, about half as fast as the TMS 9995, according to Table 2. In actual tests (Table 1) the difference is more in the 10-20% range. The mix of operations making up any particular application can produce sizable discrepancies.

With these disclaimers in mind, we can conclude that the minicomputer board pair is faster than the TMS 9995, and if we stay away from floating-point operations, the factor can be as large as 3:1 or as small as 1.5:1. Consideration must be given to the fact that this 3-times increase in speed is gained at the expense of a 125-times increase in power consumption and greater than a 10-times increase in physical size.

MIDAS Computational Configuration

In order to overcome the slowdown which would occur if we simply substituted microcomputers for the minicomputer CPU's in the MIDAS I configuration, it would be necessary to execute the microcomputers in parallel. The power of the MIDAS design stems from its thoroughly integrated designed-in ability to do exactly that. Figure 2 is a simplified sketch of a MIDAS processing segment as it is urrently operating. The eight CPU's execute identical code, and have data passed to and from them in memory blocks which are filled with preprocessed raw data from the massive data base.

In MIDAS II, the configuration might appear as in Fig. 2. Each memory has its own processor, eliminating the need for time slicing on a shared high-speed data bus. More memories (and concurrently, more processors) are needed to match the performances of the faster processors. It is normal in this configuration for processors to be idle at times; specifically during filling, emptying and during input/output-bound operations.

Time slicing in the MIDAS I configuration exacts its toll on processing speed. The memory cycle time of the minicomputer CPU is normally 200 ns. Memory operations are synchronous, resulting in this cycle time being some multiple of 200 ns regardless of memory speed. The time slicer adds 50 ns to a write cycle and 150 ns to a read cycle. The effect on the minicomputer CPU is a slowdown of data transfers by a factor of two. This effect is not as critical as it could be, because the portable memory and program memory are separate. Although accesses to data are slow, accesses to instructions are not. This overhead is more than compensated for by the use of a pipelined input processor which places discrete data events into preselected locations and by the use of an output processor which discards unfit data. The net result is that the slow-down resulting from the use of a microprocessor per-memory configuration (no time slicing) in place of the faster CPU's plus a time-slicer, will be less severe than might be expected.

Applications

MIDAS I was designed to process one particular type of data--that coming from experiments in the nuclear sciences. Consequently, its processors are required to be highly computational. Doing arithmetic, including high precision floating-point arithmetic, is very important for this application. This justifies the power consumption and size of the CPU's selected.

Another application for which the MIDAS configuration is ideal is that of performing Monte-Carlo calculations. Here again, the processing is usually compute-intensive, and more powerful processing units are a necessity. However, the pipelined input processor is not important because the data is generated within each processor and the portable memory system is useful only for output. If one of the MIDAS computational modules were to be dedicated to Monte-Carlo calculations, it may be most economical to use 120 micros, each using only 2 watts, instead of eight mini CPU's, each using 250 watts.

Sorting is a major function of processing data bases, not only in the nuclear sciences, but in most sciences relying on data collection. Sorting is also a function required of business-oriented data bases. In the sciences, sorting can often be thought of in the classical sense of comparing each element or subset of elements of data with some preset parameters and using the results of comparison to send the data into the proper bin. In business applications, generating reports is often a similar type of operation. A difference between the two is that whereas raw data from the sciences may most often be stored from a time-sequential input, business data is most often highly structured by context when it is put into the data base. Whereas it is usually practical to generate an index to data in a business data base, no such device is possible in a scientific data base. Sorting in either case is not usually a compute-intensive operation, and the micro-based MIDAS computational module would be appropriate.

Conclusions

Although the TMS 9995 is a slower device than the minicomputer CPU's being used in MIDAS I, its relative power consumption (2 watts versus 250 watts) and relative size (25 square inches versus over 250 square inches) make it an attractive alternative for use in computational modules in MIDAS II. In highly compute-intensive applications, the lack of high-speed floating point operations excludes the use of the TMS 9995.

Acknowledgment

This work was supported by the Director's Office of Energy Research, Office of High Energy and Nuclear Physics, Division of Nuclear Physics and by Nuclear Sciences of Basic Energy Sciences Program of the U.S. Department of Energy under Contract DE-AC03-76SF00098.

References

ON MIDAS:

1. Maples, Creve C., "A Specialized, Multi-User Computer Facility for the High-Speed, Interactive Processing of Experimental Data." Proceedings of Computerized Data Acquisition Systems in Particle and Nuclear Physics Conference, Santa Fe, NM, May 14-17, 1979.

2. Maples, Creve., Proposal for a High Speed Interactive Facility for the Reduction and Analysis of Scientific Data. Presented at Asilomar, CA, meeting of the American Physical Society, Nov. 1-3, 1978. LBL-7196 (1978), Lawrence Berkeley Laboratory, Berkeley, CA 94720.

References 3-5 presented at Topical Conference on Computerized Data Acquisition in Particle and Nuclear Physics, Oak Ridge, TN, May 28-30, 1981. All appear in Proceedings of the conference.

3. Maples, C., Rathbun, W., Meng, J., and Weaver, D., "A Fast Time-Sliced Multiple Data Bus Structure for Overlapping Data Transfers and Transformations.

4. Maples, C., Rathbun, W., Weaver, D., and Meng, J., "The Design of MIDAS - A Modular Interactive Data Analysis System."

5. Maples, C., Weaver, D., Rathbun, W., and Meng, J., "The Utilization of Parallel Processors in a Data Analysis Environment."

6. Meng, J., "Power Basic and the 9980/9981. TIMIX 1981 National Symposium, New Orleans, LA, March 8-11, 1981. LBL-12235, Lawrence Berkeley Laboratory, Berkeley, CA 94720.

7. Meng, J. and Weaver, D., "Use of Embedded Microcomputers in System Debugging and Maintenance," Proceedings of the ICEI 1981 Int. Conference, IEEE Applications of Mini and Microcomputers, pp. 297-301.

TABLE 1

Power BASIC Loop	990/101 Execution Time (Sec)	TMS 9995 Execution Time (Sec)	Improvement(%)
10 FOR X = 0 TO 1000 20 A = SIN (.5) 30 NEXT X 40 PRINT A 50 GOTO 10	17	15.3	10
5 B = 6.1::C = 9.23::A = 0 10 FOR X = 0 TO 1000 20 A = B + C + A 30 NEXT X 40 PRINT A 50 GOTO 10	4	3.3	17.5
10 FOR X = 0 20 A = CRF(0) 30 NEXT X 40 PRINT A 50 GOTO 10	22	19	13.6

In spite of its 8-bit data bus, the TMS 9995 outruns the TMS 9900 by 10-20 while executing Power Basic. The TMS 9900 was run without memory wait states. The TMS 9995 was run with one wait state per memory request.

TABLE 2

	MINI CPU/EAU	TMS 9995	Ratio (9995/MINI)	990-101	Ratio (9900/9995)
Physical Size	2 cards 14" x 19-1/2" No memory	5" x 5"	1/10.9		
Power Consumption	50 amp 5 V (250 W)	~.4 A 5 V (2 W)	1/125	~20 W	
Load Reg. Fr. Mem.	1.1 μs	2.5 μs	2.3	7.3 μs	3.2
Branch	1.1 μs	1.7 μs	1.5	3.6 μs	2.4
Load to Reg. Add fr Mem Store in Mem	3.3 μs	4.3 μs	1.3	7.3 μs	1.7
Shift	.5 - .7 μs	2.7 - 8 μs	5.4 - 11.4	4 - 15 μs	1.5
Sing. Prec. Divide	6.1 μs	~ 15 μs	~ 2.5	~31 - 42 μs	~ 2 - 3
Sing. Prec. Multiply	2.7 μs	~ 9 μs	~ 3.3	~ 18 μs	~ 2
Floating Point Arithmetic	1.6 - 14.4 μs	N/A		N/A	
System Clock	15 MHz	3 MHz	1/5	3 MHz	1

For most operations, the minicomputer CPU outruns the TMS 9995 by about a factor of two. Shifting and arithmetic operations widen the difference considerably. There is no floating point hardware on the TMS 9995 for comparison. For many types of operations, the 125-times increase in power consumption and the 11-times increase in board area improved performance by less than a factor of 2.

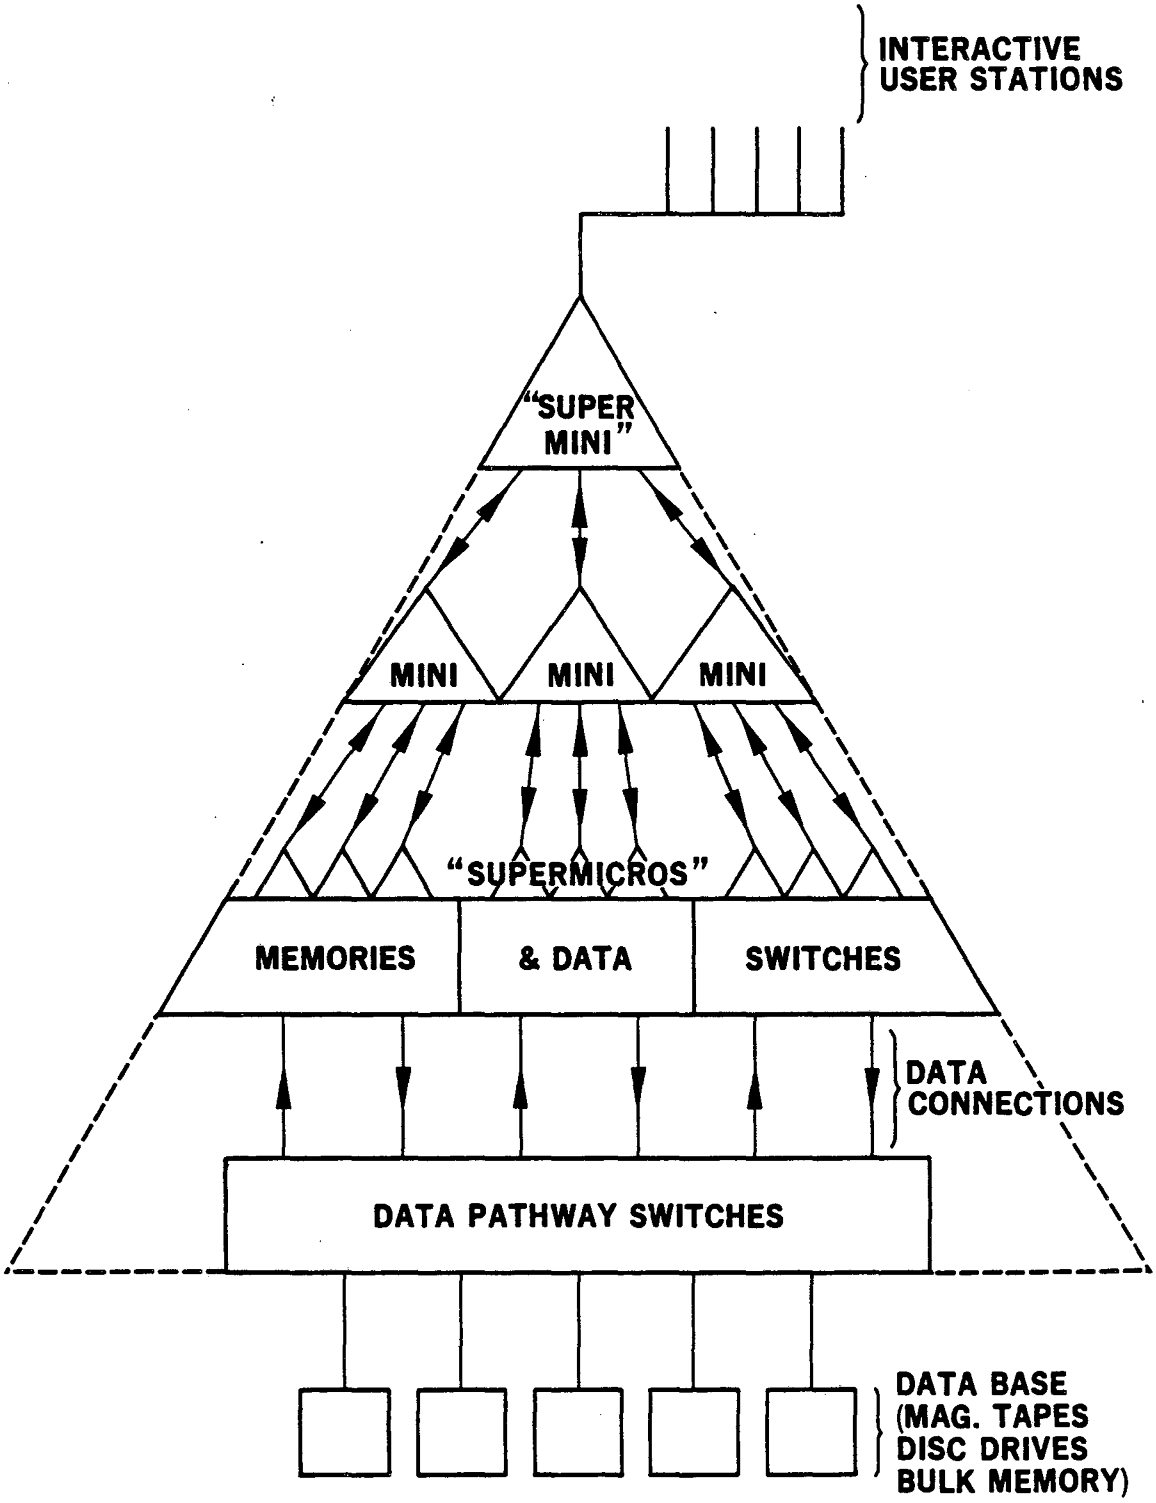

XBL 817-10749

FIGURE 1. MIDAS - Computational modules reside at the "supermicro" level.

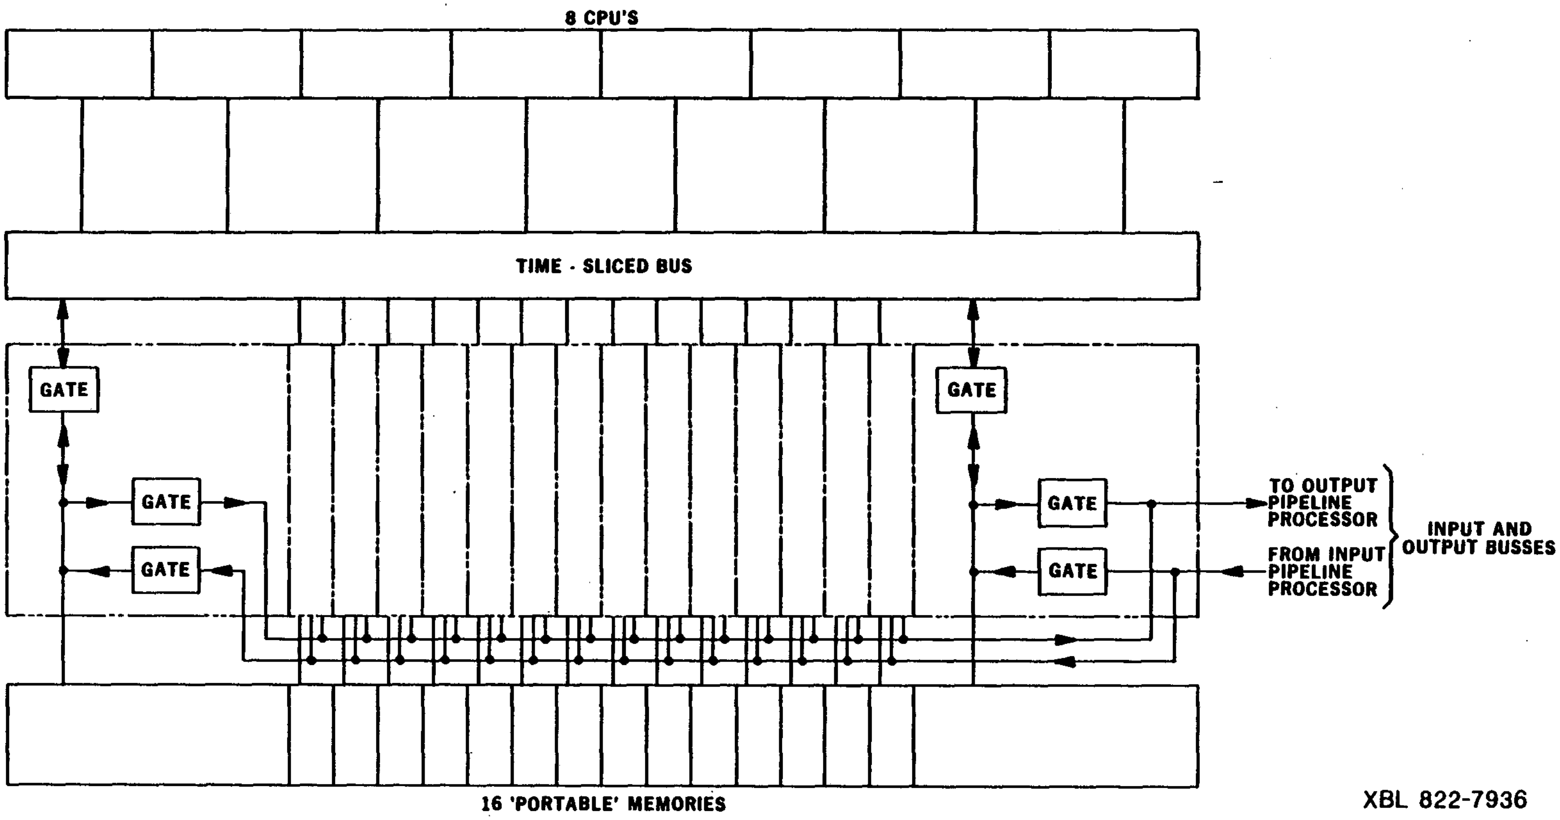

XBL 822-7936

FIGURE 2. Close-up simplified view of MIDAS processor array. Eight CPU's share sixteen 'rotating' memories through a single time-sliced bus.

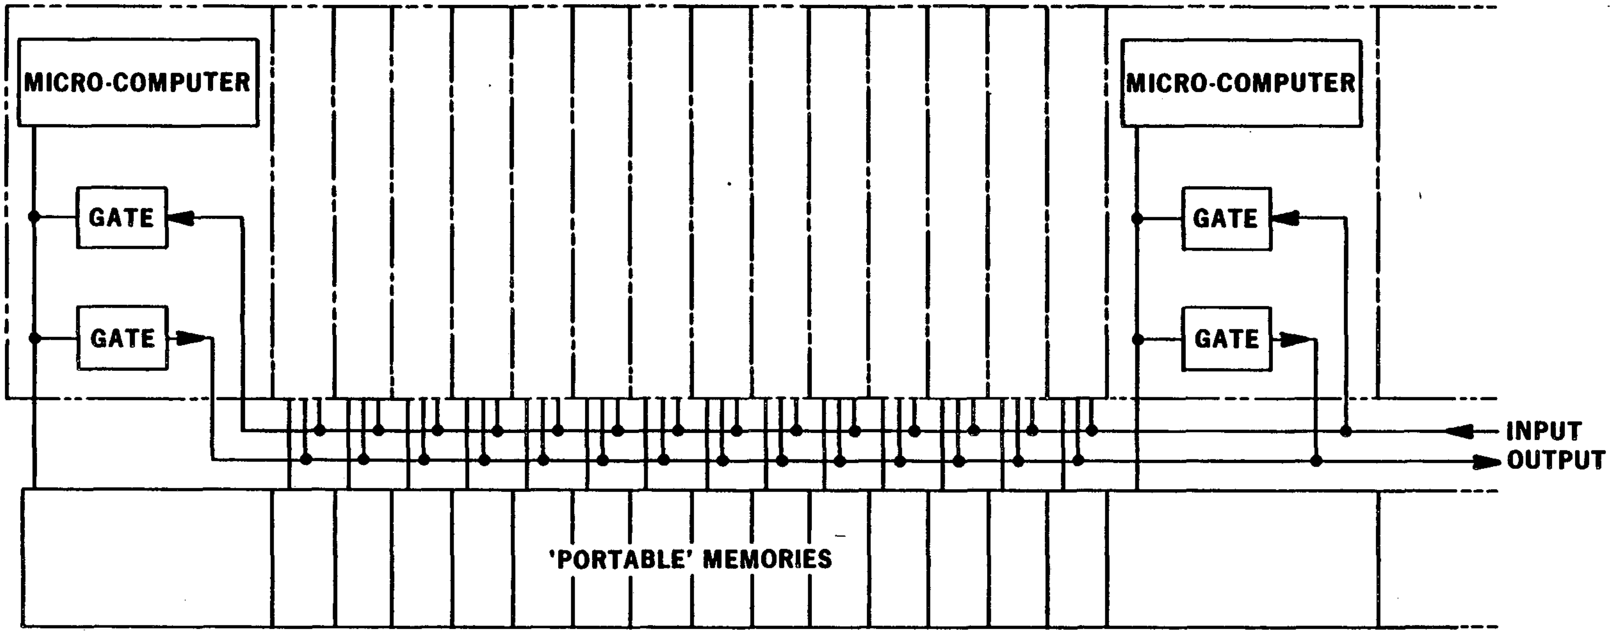

XBL 822-7937

FIGURE 3. In this proposed configuration the time-sliced bus is eliminated and a CPU (micro-computer) is attached to each 'portable' memory. More memories are needed to enable the overlap of filling and emptying with processing.

This report was done with support from the Department of Energy. Any conclusions or opinions expressed in this report represent solely those of the author(s) and not necessarily those of The Regents of the University of California, the Lawrence Berkeley Laboratory or the Department of Energy.

Reference to a company or product name does not imply approval or recommendation of the product by the University of California or the U.S. Department of Energy to the exclusion of others that may be suitable.

TECHNICAL INFORMATION DEPARTMENT
LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA
BERKELEY, CALIFORNIA 94720