Unconventional systems*

by DANIEL L. SLOTNICK
Department of Computer Science
University of Illinois
Urbana, Illinois 61803

INTRODUCTION

Background observations

For the bulk of the past decade, the computer manufacturer* has maintained the "party line" that their successive models of high-end machines possessed all the performance that would be required for the foreseeable future. (This "party line" alternated with abortive attempts to, in fact, build high performance, compatible extensions of the product line which turned out to be neither high performance nor compatible.) While it appears to be true that currently available equipment will for a while continue to be adequate to permit plumbing supply houses to do inventory control, it has always been outrageously false that this equipment come evens close to meeting our scientific or military needs. In these areas, we have witnessed the oft noted paradox of technological advances spurring scientific needs in such a fashion that the disparity between what is needed and what is available continues to grow. In the case at hand, this is no paradox at all in that it is only now that we can realistically envision computers of sufficient speed and capability to make it possible to start using them to solve problems with a non-trivial intellectual content.

Two related sequences of events have in very recent years made it no longer necessary to apologize for grown men to be concerned with producing computing machines of vastly greater capability. These are the consequential gains made by a technology (no thanks whatever to the manufacturer) that was not obliging enough to stand still and secondly that we have the possiblity of performing for the first time with the latest generation of equipment, calculations that couple closely to the mental and physical processes of scientific advancements.

Current status review

A startling result of the past decade of progress in switching elements is the number of such elements which can be employed in a system of given reliability. For example: A system which contains 10 to 20 thousand vacuum tubes achieved mean-time between failures of at most several tens of hours. Comparable or even somewhat greater mean-time between failures are achieved today by systems with upwards of 10^6 conventional transistors. Within the next five years with the coming large scale integration, this number should go to between 10^8 and 10^9 discrete transistor equivalents.

Now, granting the assumption that we will in the foreseeable future be limited by computer capacity in a broad spectrum of areas, it is legitimate to ask that organizational concept be evolved which can utilize a greater number of components to achieve at least proportionally greater performance. It is an important practical aside to note that during the same interval of time the cost of a small signal switching element has gone down to an extent permitting the employment of the maximum number which in non-redundant designs yield a reliable system.

Another important factor affecting our approach to computer design is the growing trend to employ design mechanization. Although this field is still very young it has already advanced sufficiently for us to forecast confidently that reduced development costs and lead times should continue to make successively bolder progress tolerable. I am by no means suggesting that computers of vastly greater power will be designed more cheaply and more quickly than their predecessors, but that as a percentage of our gross national product, these costs will be tolerable in their relative magnitude and continue to yield the nation a reasonable return.

*The term manufacturer is used here only to describe large manufacturers, say those possessing more than half of the total computer market.

The manufacturer continues to bombard our (mercifully) somewhat desensitized ears with the problem of software conversion in an effort to perpetuate hardware that when looked at even ten years from now, will have value only to the historian or antique collector. (In the case of recent compatible families this argument has changed somewhat in that the manufacturer finds that offering even paltry improvements in equipment eliminates compatibility. The new argument is that only the manufacturer can produce software in a timely, economical fashion because of his ability to pool requirements from different customers. The supreme irony of this point of view is already legion.) The value of the machine software investment (notice I said value not cost) has been greatly exaggerated for the same reason by the manufacturer. In particular, the sophisticated user (i.e., the academic and military eggheads with their huge problems) have traditionally been required to program their problems in a low level language. The argument that we must make no further progress in computer organization—a field that is no more than 20 years old in order to maintain a fictitious compatibility with Baggage's computer and its electronic progeny is arrogant and patently absurd.

In seeking up-to-date criteria for judging the efficiency of organizational concept the following are obviously suggested:

(1) That they make efficient use of high-level semi-conductor components (i.e., equivalent to from several hundred to several thousand discrete elements), and

(2) That the designs keep a maximum number of memory amplifiers working, i.e., maintain maximal traffic on memory distribution busses both productively and without chaos.

A third criterion is the extent to which a given design can employ design mechanization in its implementation. An extention of this last criterion threatens at some time in the future to explode all our conceptions about computing machines. It is already possible to conceive a machine which will automatically design and fabricate special purpose computers on the basis of a set of instructions not differing greatly from a current machine program. In other words, the filing cabinets full of punched cards would be replaced by filing cabinets full of special purpose computers. It is my current plan to pursue such machines when the current machine we are building at the University of Illinois (ILLIAC IV) is successfully operating. I am greatly heartened that some of the same people who five years ago reacted to the idea of ILLIAC IV with smug incredulity are reacting now in the very same way to this idea.

Some approaches to large computing capability

To return to 1967: What are the approaches currently being explored to achieve greater speed? The manufacturer has a simple answer. If a problem requires a 1000 times the power of computer "X," use a 1000 computer "X's." Of course, this ignores the fact that attempts to achieve cooperative interactions between general purpose computers using clip leads and other strong medicine have in the past yielded lamentable and sometimes laughable results. This approach we can dismiss out of hand if for no other reason than that it is ponderously inelegant; but there is another reason—it doesn't work.

Another approach is to simply state that the conventionally organized general purpose computer will continue to improve so as to keep pace with requirements. This argument ignores the fact that the requirements which are not generated by the existence of a given computer at a given place, for example those motivated by scientific or military problems where the outside world determines the problems, have already produced requirements greatly in excess of current capabilities. Secondly, as has been pointed out now for many years, improvements in general purposes machines are not open-ended because of speed of light and memory access and buss distribution limitations.

Next, in the current set of ideas to achieve high performance are a number of concepts which achieve their very high performance partially by sacrificing some generality of purpose. I must remark here about the erroneous impression created by referring to our current machines by the name "general purpose" computer. This name makes the false implication that they are uniformly good (or poor) at all things. Moreover, the only high performance computer* currently in operation differs markedly in its organization from its predecessors and moreover experience with this machine seems to indicate that the variation in its efficiency of application from problem to problem is quite considerable.

The "pipe line" systems to be available in the next several years offer another order of magnitude of performance over the fastest machine currently available.

These "pipe line" machines achieve speed through the same basic mechanism as parallel machines, i.e., applying a single instruction to a considerable block of data. They, as a consequence, suffer from the same general disadvantage, namely: If the problem does not permit the data to be treated in blocks, their efficiency degrades. Most problems examined, how-

---

*This is of course, the Control Data 6600 soon to be succeeded by a related machine which is four times faster.

ever, that require very high performance do seem to possess this property of permitting data to be handled in blocks to a remarkable degree although this is not always obvious at first glance. The "pipe line" approach is novel and merits our closest attention. It is being pursued by the major manufacturers in the field.

The approach that I am personally pursuing is the parallel approach. It is the only current approach which is open-ended. The same remark holds for both parallel and "pipe line" systems, namely: Its generality is not fully established.

I will furnish a brief overall description of the system here and refer you for details to a separate session at this conference dealing exclusively with the ILLIAC IV hardware, software and applications.

The ILLIAC IV system

Simply to illustrate the capacity which current technology permits in a computing system, I will give a brief description of the ILLIAC IV system being developed jointly by the University of Illinois and industry.
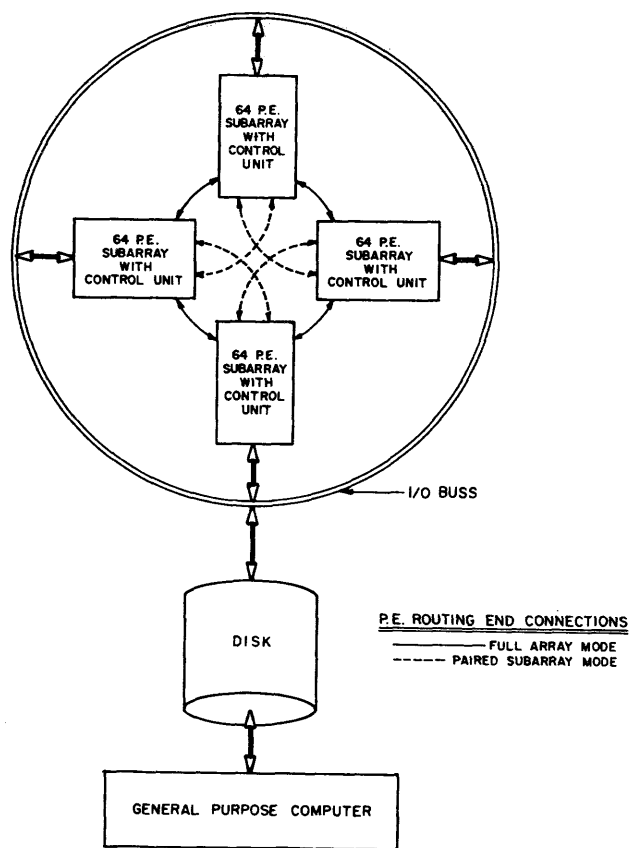

Figure 1 – ILLIAC IV General Organization

Figure 1 illustrates the general arrangement of the ILLIAC IV system, centered on four subarrays, each containing 64 processing elements under the direct control of a subarray (Figure 2) control unit.

Data is transferred between the memory units of the processing elements of the subarrays and a large scale disk file buffer memory via a highly parallel input/output buss. Such input/output transfers are controlled by an external general purpose computer which also supervises the ILLIAC IV program runs. The general purpose computer is provided with a limited set of connections to the subarray control units for this purpose.
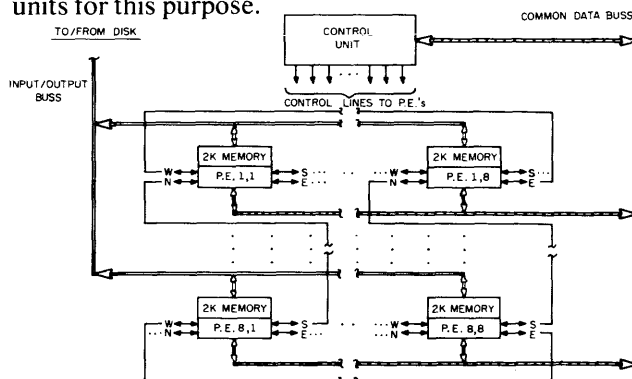

Figure 2 – Organization of 64 P.E. Subarray

The routing connections of the processing elements in the four subarrays are arranged to permit a *united mode* of operation, in which the four subarrays act as a single large array for routing purposes, a *paired subarray mode* of operation which provides routing for two arrays, each containing two of the subarrays, or an *uncoupled mode* in which each subarray operates independently.

Figure 3 shows the structure of one Processing Element (PE). Each PE is provided with three 64-bit arithmetic registers and high speed adders for full 64-bit floating and fixed point operations. The processor is also provided with 2000 64-bit words of thin film memory.

Four routing connections, identified as North, East, South and West are provided in the subarray as shown in the figure together with busses to and from the control unit and the disk unit for global and input/output data.

It is a vital consideration that the PE is regarded as a complement. It is, in fact, delivered in assembled, tested form to the system supplier by the semiconductor manufacturer. Provision is made in the design of the PE to permit an orderly transition from hybrid LSI components to monolithic LSI components during the course of the program. The technology employed in the design and fabrication of the PE will, it is anticipated, yield a unit price for the first system of this very powerful 64-bit floating point unit of under $10,000.

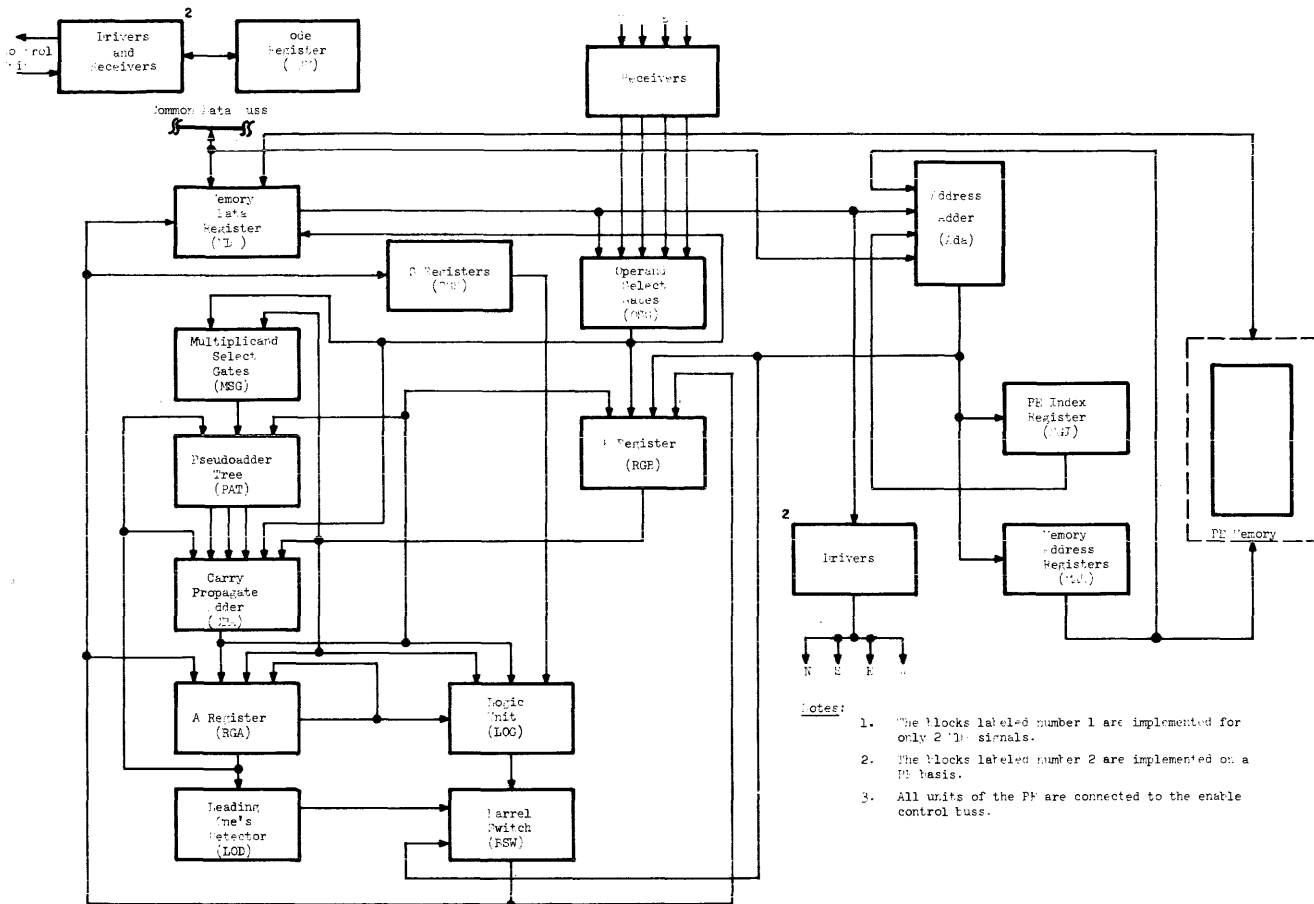Detailed analytical and programming work has been done in the following areas:

Figure 3 — Processing Element Component, Block Diagram

(1) Weapons Effects
   a. Eulerian flow
   a. Eulerian flow
   b. Lagrangian flow
   c. Neutron transport
   d. Underground stress-strain model
(2) Alternating Direction Implicit Relaxation
(3) General Circulation Weather Model
(4) Matrix Operations
   a. Inversion
   b. Eigenvalue calculation
(5) Linear Programming
(6) Multichannel Filter Design and Convolution
(7) Fourier Transformation

In all the areas listed it was found that highly efficient methods could be evolved, i.e., methods which keep almost all PE's running almost all the time. Thus, the speed-ups observed have been in the order of 256 times N where N relates the speed of the PE to the speed of the computer selected for comparison.

As an illustration of the sort of consideration involved in securing high efficiency on a parallel machine, Figure 4 shows the memory allocation for a 5×5 matrix in 5 PE memories (PEM's). We refer to this storage scheme for matrix elements as skewed storage.

The main reason for skewing a matrix is that rows and columns of a matrix can be accessed with appropriate PE indexing. For instance, to fetch elements
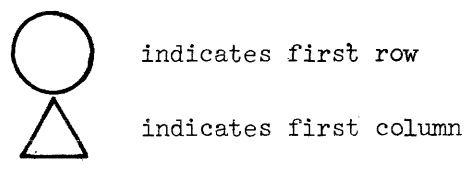

Figure 4 — Skewed Storage Technique for Matrices

of a column, neighboring PE's access elements differing in address by 1.

The illustrated matrix fits nicely into the memory allotted. For matrices where this is not true packing schemes exist to minimize wasted storage space in relation to the operation being performed.

The case of skewed storage just described is an important but simple example of the unique considerations made necessary by the parallel organization of this computer. Manual calculation (like man himself) is essentially sequential. Our electronic calculating aids (computers) have hewn strictly to this "natural" organization of calculations. It is, however, by no means "natural" for a technological innovation to so restrict itself. In fact, man is a slow, highly error prone computing device. The parallel approach to computing, it must be said however, does require that some original thinking be done about numerical analysis and data management in order to secure efficient use. In an environment which has represented the absence of the need to think as the highest virtue this is a decided disadvantage.