RADCAP—An operational parallel processing facility

by JAMES D. FELDMAN and LOUIS C. FULMER

Goodyear Aerospace Corporation
Akron, Ohio

SUMMARY

An overview is presented of RADCAP, the operational associative array processor (AP) facility installed at Rome Air Development Center (RADC). Basically, this facility consists of a Goodyear Aerospace STARAN* associative array (parallel) processor and various peripheral devices, all interfaced with a Honeywell Information Systems (HIS) 645 sequential computer, which runs under the Multics time-shared operating system. The RADCAP hardware and software are described only briefly here since they are detailed in companion papers presented at this conference.[1,2] The latter part of this paper dwells on the objectives of the RADCAP facility and plans for its use.

The STARAN associative parallel processor is a processor based on an associative or content addressable memory and a related ensemble of bit serial processing elements. STARAN is considered to be the first practical associative processor ever produced.[3] This claim of practicality is based on the fact that the design concept for the associative memory of STARAN allows the use of the same high-volume, standard, large-scale integrated (LSI) circuit memory devices that are in widespread use by the computer industry. In fact, every electronic component used in the STARAN associative parallel processor is available from your local components distributor. The significance of this fact is that now, for the first time, associative processors enjoy the same cost per bit of storage as does the conventional computer.

HISTORICAL BACKGROUND

From the time Slade and McMahon first described their catalog memory[4] in 1957, many attempts have been made to implement an associative memory. Some of these attempts were successful, but until recently none has been very practical. Table I lists some of the device technologies that have been used in the past to implement associative memories in the laboratories and in a few experimental models. Except for a few special applications of plated wire, none of these device technologies have ever been used successfully for conventional memory technology. Further, for associative memory applications, none of these device technologies have been very practical. All had one common characteristic—high cost per bit of storage. The exotic nature of the device, the custom nature of the associative cell, and the resulting low volumes were the factors contributing to the high cost per bit.

The extension of an associative memory to an associative processor by the addition of serial arithmetic units to each word of associative storage was demonstrated using plated-wire technology and was reported[5] in 1970. At about that same time a new version of the STARAN associative parallel processor was in the formative design stages at Goodyear Aerospace Corporation. Two choices for the device technology of the associative array were available: plated wire or LSI. Plated wire was quickly discarded as a "squeezed" memory technology caught between the well established magnetic core memory and the emerging solid-state integrated circuit memories. The choice was to go with LSI, and the temptation to design a custom LSI associative array was great.

The design of the associative array could be made parallel by word and bit and could include a repetitive cellular structure that would lend itself nicely to an LSI array. The cell could be designed to include all of the desirable characteristics of an associative processor: nondestructive readout storage, logic for the associative and arithmetic functions, and access to the cell in both the word and bit slice directions.

Figure 1 is a simplified diagram of a two-dimensional custom LSI associative array, where each cell has been designed to include the desirable storage, logic, and access characteristics of an associative processor. The largest neg-

TABLE I—Associative Memory Element Technologies

Cryogenics
Cryoelectrics
Multiaperture Ferrites (MADS, MALE, BIAX)
Ferroelectrics
Toroidal Cores (FLUXLOK, BILOC)
Discrete Transistor Associative Cells
Discrete Integrated Circuit Associative Cells
Plated Wire
Custom LSI

* TM, Goodyear Aerospace Corporation, Akron, Ohio.

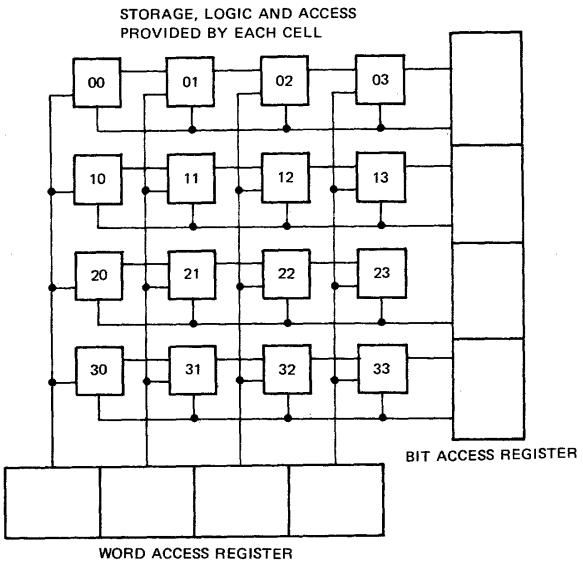

Figure 1—Custom LSI associative memory (parallel by word and bit)

ative technical factor in this approach is the large number of gates for each associative cell. In one design[6] of this type, the associative cell required about 40 gates. The high gate count implies more silicon per cell and, therefore, lower yield and higher cost. A second technical factor is the large number of pins required to package the two-dimensional LSI array. Another design[7] that packages 128 associative cells requires a 40-pin package and 256 cells requires a 56-pin package. Tradeoffs between pin count and logic complexity are possible but limited. The high pin count results in the use of expensive nonstandard integrated circuit packages. A third technical factor in a design of this type is that heat dissipation, geometry, and economics dictate the use of external sense amplifiers.

Even if all of the above technical problems could be solved economically, the largest problem is a nontechnical one—that of low volume. Due to the custom nature of the device, it can only be used in the associative processor for which it is designed, and unless an integrated circuit is produced in response to a tremendous volume demand, its cost will always be relatively high. Related to this custom device problem is the fact that the designer of the associative processor must bear the high nonrecurring costs by himself, usually with a single source which further insures relatively high prices, and without the benefit of a background of reliable data for the specific device. Then, due to the high nonrecurring investment and the low volumes, the associative processor designer will not be able easily to take advantage of the technological advances that are occurring rapidly in the integrated circuit industry and he soon has an obsolete device on his hands.

Further, the integrated circuit industry does not want to be involved in low volume production programs. An excerpt from a 1973 report for the government entitled "Approaches to Custom Large Scale Integration"[8] is quoted below to support this point.

"It is increasingly apparent that LSI offers considerable benefits, most of them related to cost, in the implementation of the digital portion of any system of reasonable complexity, provided it is produced in high volume. High volume can be taken for granted in standard LSI devices, such as memories. Such is not the case for logic, however, which, in order to be near optimum, must be custom for each function; the volume in which it is produced therefore depends entirely on the number of systems to be built. While the total number of systems required by the military may be large, the number of any one system type is often small. However, in order to reap the full cost benefits of LSI, the volume must be high, at least of the order of $100,000 to $200,000 per year per chip, and preferably higher. Only at this level does it become economically feasible to fully optimize the design by handcrafting for maximum area utilization and performance, to "tweak" the process for maximum yield, and to package at the lowest cost. These cost savings are not realized at lower volume.

The following summarizes the major comments made by the industry:

- *High volume is the most important business criterion.*
- *Low-volume LSI development work was done principally with military funding. Work stopped when funding dried up.*
- *Low-volume LSI is bad business for the big semiconductor companies.*
- *Only systems houses with captive IC capability can respond to low-volume LSI requirements.*
- *Manufacturing in IC houses is geared to very large lots.*
- *The customer must be very sophisticated if he wants low-volume LSI, preferably he generates his own design.*
- *Limiting resources are: 1. Capital equipment; 2. Manpower. These cannot be wasted on low-volume businesses.*
- *Every effort should be made to use standard products.*

The custom LSI approach was considered for STARAN but was discarded primarily due to the nontechnical factors discussed above. A custom LSI approach, although technically appealing, would result in an impractical implementation of the associative processor. That approach can never break out of the "cost-volume hangup."

The GAC approach to using LSI for the associative processor was just the opposite of the custom LSI approach. Instead of combining the requirements of storage, logic, and access into one custom LSI chip, the requirements were divided to see if they could be implemented using several standard, high-volume large-, medium-, and small-scale integrated circuit chips. Figure 2 is a schematic representation of the approach to the problem. If this problem could be solved, the resulting cost of associative processors would be low and the "cost/volume hangup" would be solved.

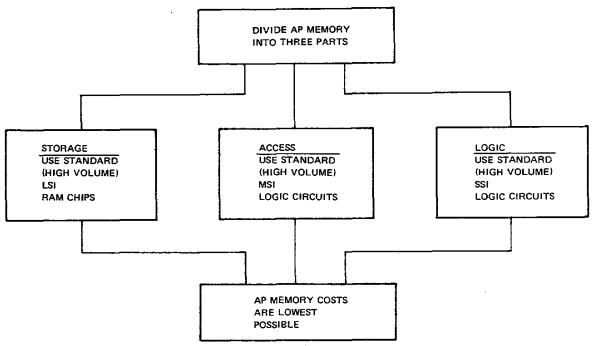

Figure 2—Approach to implementing an associative processor memory

Storage

Figure 3 shows a typical standard LSI 1024×1 memory chip in a 16-pin package. These memories are (or soon will be) available in high-volume production in various device technologies (MOS, TTL, ECL, CMOS) and are in widespread use in the conventional computer industry for mainframe memory. Since they are committed to replace the magnetic core being used for that purpose, the volume leverage for lower cost is already at work. For use in a serial by bit (slice) associative processor, this memory chip is functionally equivalent to single plated wire and can satisfy the nondestructive storage requirement very nicely. The number of gates per bit of storage is approximately two.

Logic

In a serial by bit associative processor, the logic necessary to perform the associative and arithmetic functions is embodied in a small bit-serial processing element sometimes called a response store or a serial arithmetic unit. Typically, this processing element has a complexity of about 32 gates and consists of three or four flip-flops and some logic gates.

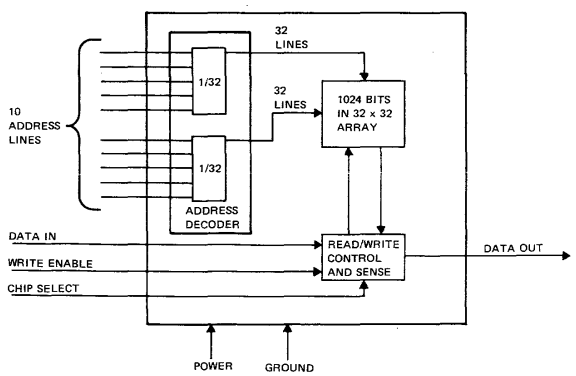

Figure 3—Standard LSI random access memory (1024-bit chip in 16-pin package)

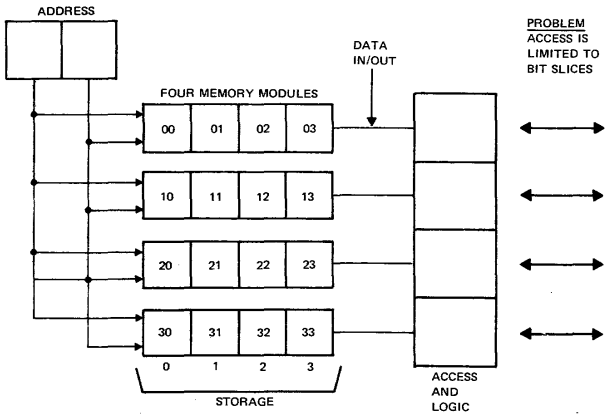

Figure 4—Standard LSI memory used in an associative processor

It is easily constructed of small-scale integrated (SSI) circuits. When operated with a 256-bit store, the processing element adds 32/256 or 1/8 gate per bit of storage. With a 1024-bit store, the ratio is 32/1024 or 1/32 gate per bit of storage.

Access

Figure 4 shows a simplified diagram of four LSI memory chips connected to four serial processing elements. This organization is very similar to the plated-wire serial by bit associative processor mentioned earlier. The problem with this design is that access to a bit slice is accomplished in one-bit read or write time but access to a word or part of a word requires (n) bit read or write times (where n is the number of bits to be read or written). It would be desirable to access either a bit slice or a word slice in one read or write time. There may also be cases where it is desirable to access in a mixed mode (words and bits) addressing technique that allows up to 256 cells to be accessed at once.

The problem of access to either bit slices or word slices or combinations of the two has been solved with a proprietary GAC design that uses a logical network between the memory and the processing elements. This network is called a flip

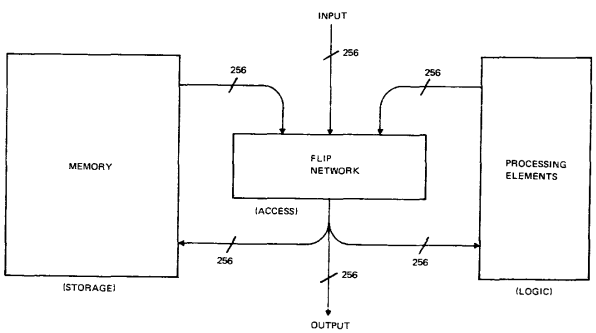

Figure 5—Organization of a STARAN associative array

RADCAP FACILITY

Figure 8 shows a block diagram of the hardware within the RADCAP facility. The 645, which has been in existence at RADC for several years, is a very large computer system with a multitude of peripherals typical of large time-shared systems. In March 1973, hardware was delivered to RADC in the form of a STARAN parallel processor with four arrays, a custom input/output unit (CIOU), a hardware performance monitor, and a variety of peripherals. Subsequently, the CIOU was used to interface STARAN with a 645 I/O channel. At the same time, STARAN software was interfaced with the 645 Multics time-shared operating system.

At present, the RADCAP facility is totally operational and includes system software to allow for operation in both a STARAN stand-alone mode and an integrated STARAN/Multics mode.

STARAN PARALLEL PROCESSOR

STARAN can perform search, arithmetic, and logical operations simultaneously on either all or selected words of its memory. Figure 9 shows the basic STARAN elements. The most important are the associative array and its unique multi-dimensional access capability, which, along with the other elements, are described in more detail in referenced publications.[1,3,9] Listed below are brief descriptions of the STARAN elements:

1. Associative array: provides multi-dimensional access, content-addressable memory with 65,536 (2^{16}) bits of storage and 256 processing elements; permits parallel arithmetic, search, and logical operations.
2. AP control: performs data manipulation within associative arrays as directed by program stored in AP control memory.
3. AP control memory: stores AP control instructions. Can also store data and act as buffer between AP control and other system elements.
4. Sequential controller and memory: performs maintenance and test functions, controls peripherals, maintains job control, provides means for operator communication between various STARAN elements, and assembles STARAN programs written in MAPPLE (*M*acro-*A*ssociative *P*rocessor *P*rogramming Langua*g*e*E*).
5. External functions: transfers control information among STARAN elements.

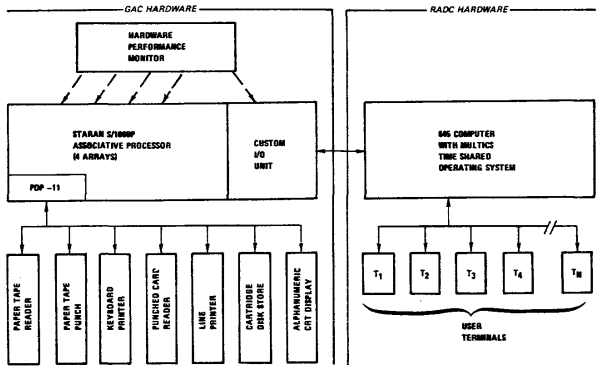

Figure 8—RADCAP facility

STARAN has been designed to provide a flexible I/O capability. The standard peripherals for STARAN are listed below, along with a typical list of optional peripherals:

1. Standard: cartridge disk drive and control, paper tape reader, paper tape punch, and keyboard printer.

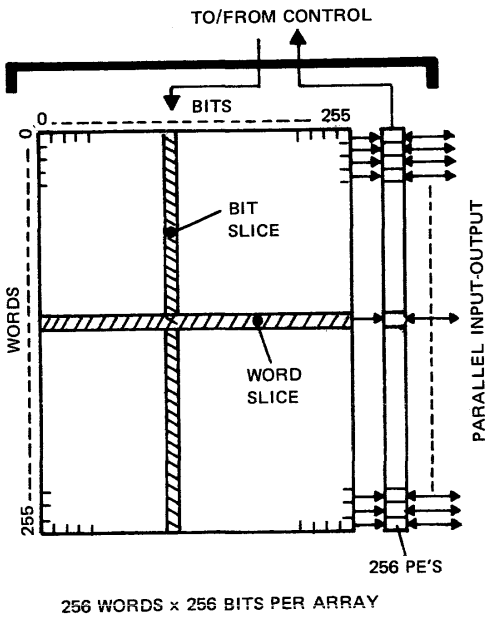

Figure 7—STARAN associative processor array

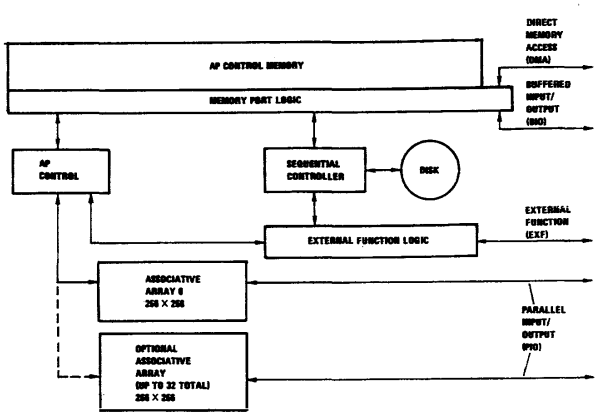

Figure 9—STARAN block diagram

2. Optional: line printer, card reader, magnetic tape, keyboard crt, and other peripherals, as desired, that are compatible with the Digital Equipment Corporation (DEC) PDP-11.

All these peripherals interface with the STARAN system's sequential controller, a PDP-11 mini-computer. STARAN also provides facilities for interfacing with other processors. The four buses provided (see STARAN block diagram, Figure 9) are the direct memory access, the buffered I/O, external function, and parallel I/O.

The direct memory access is a 32-bit bus for STARAN to address external memory. The AP control or the sequential controller can access external memory at a rate dependent upon this memory's cycle time.

The buffered I/O is a 32-bit bus for processors to address STARAN. Depending upon which portion of control memory is accessed, the access rate is 0.4 to 1.0 microsec per 32-bit word.

The external function is a bus for exchange of control signals. Discrete signals and interrupts can be both generated and accepted across this bus.

The parallel I/O is a bus for STARAN array I/O. Up to 256 bits per array (e.g., one bit per array word) can be provided. If all 32 arrays are implemented, up to 8192 bits can be utilized in parallel at a transfer rate less than one microsecond, dependent upon the desired application.

STARAN PERFORMANCE SUMMARY

In a high-speed, asynchronous, pipe-line type processor such as STARAN, it is difficult to summarize performance since speeds vary with instruction types, types of loops, etc. Also, the overall effective speed depends upon the number of words in the arrays over which the simultaneous operations are occurring. However, an effort is made below to list the performance and the features of 256×256 associative array, the control unit, and the interface portion of STARAN:

Associative Array Features

- Up to 32 arrays per system
- Multi-dimensional access (bit slice or word slice)
- Array module speed:

Typical search:	150 nsec/bit
Typical add or subtract:	800 nsec/bit
Read bit or word slice (256 bits):	150 nsec
Write bit or word slice (256 bits):	300 nsec

Control Unit Features

- Two separate processors: AP control, sequential controller
- Solid-state control memory capacity: 2K×32 standard, 4K×32 maximum
- Solid-state control memory speed: 150 nsec/instruction (typical)
- Bulk core capability: 16K×32 standard, 32K×32 maximum
- Bulk core speed: 1 microsec (read or write)

Interface Capabilities

- STARAN to address external memory: rate-memory dependent
- External processor to address STARAN: 0.4 to 1.0 microsec/32-bit word
- Parallel I/O to/from associative arrays: less than 1.0 microsec/8192 bits (maximum)
- Control signals and interrupts

Custom Input/Output Unit (CIOU)

Figure 10 shows a simplified block diagram of the STARAN/RADCAP custom input/output unit (CIOU). As indicated, the CIOU contains a parallel input/output (PIO) module, a 645 computer interface, and an internal performance monitor. The CIOU functions as a mini-processor much the same as the control unit portion of STARAN. Processing within one array module (e.g., under STARAN control) may be concurrent with I/O in another array module (e.g., under PIO control).

As directed by instructions stored in PIO control memory, the optional PIO module manipulates data among and within the associative arrays concurrent with operations as directed by AP control. The PIO module contains eight ports, with 256 bits per port to accommodate associative array I/O and to permute data.

The 645 interface logic provides a communication path between the 645 computer and the STARAN system. This interface logic contains a 30-character queue and a 32-bit status register, which are tied to a 645 I/O channel. The status register contains interface control signals, and the queue buffers data being transferred to or from the 645.

The internal performance monitor, although contained in the CIOU, is best discussed in the following description of the hardware performance monitor.

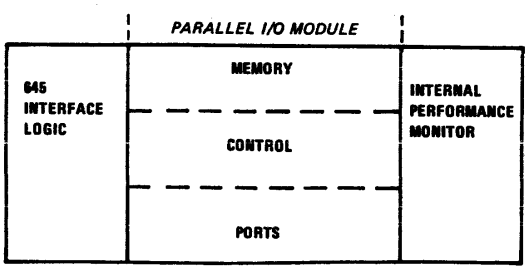

Figure 10—Simplified block diagram of custom I/O unit

Hardware performance monitor

To help meet a RADCAP facility objective of measuring system performance, a hardware performance monitoring capability has been provided by an internal performance monitor in the CIOU cabinet and an external performance monitor system. Measurements can be made to determine instruction execution timing, control memory and bus utilization, array utilization, and activity in the pager, the PIO module, and the 645 interface.

The internal performance monitor is used exclusively for STARAN instruction execution times and instruction event times. The events counted and timed are the execution of flagged instructions in AP control. Between a start flag and an end flag, a timer increments at a 100-nsec rate. Overflows from this counter interrupt the sequential controller. In addition, the sequential controller can interrogate the event counter and timer.

The external performance monitor is a self-contained system that can monitor any point of STARAN or the custom I/O. Data are acquired via probes that detect logical signal changes in either an event count or elapsed time mode. Several probes can be logically connected via a patchboard to trigger a counter. At regular intervals, the contents of the counters are written as a record on a magnetic tape unit. The performance monitor software then evaluates the collected data and produces the results in the form of reports and graphs. The software for the performance monitor runs on the 645.

Physical description of hardware

All the elements shown in the STARAN block diagram (Figure 9), including the associative arrays, are built using dual-in-line IC's (integrated circuits) mounted on multilayer printed circuit boards. Thus, the physical construction of STARAN and the CIOU is similar to that of typical high-speed sequential processors.

Figure 11 shows Goodyear Aerospace's STARAN demonstration and evaluation facility. Table III gives the approximate numbers of cabinets, boards, and IC's for the various STARAN models. These figures do not account for I/O logic, since this varies from one installation to another. The STARAN/RADCAP CIOU, which includes the parallel I/O option for all four arrays, contains approximately 200 boards and 8,000 IC's.

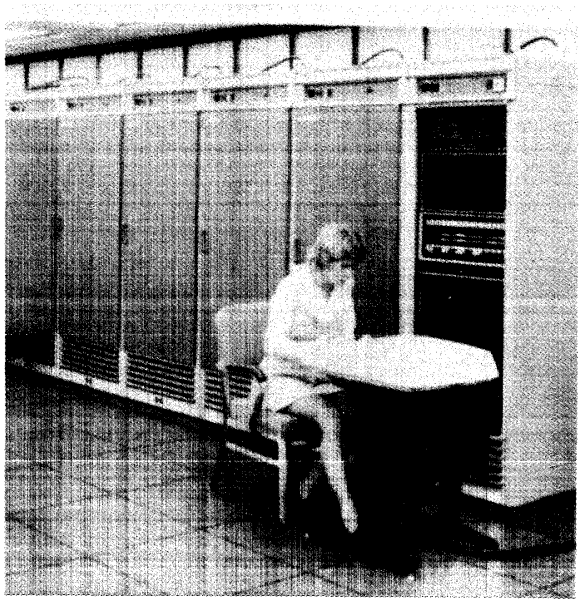

Figure 11—STARAN demonstration and evaluation facility

TABLE III—Approximate STARAN Component Count*

STARAN* Model	No. of Arrays	No. of Cabinets	No. of Printed Circuit Boards	No. of Integrated Circuits
S-250	1	3	220	9,000
S-500	2	3	276	11,500
S-750	3	3	332	14,100
S-1000	4	4	412	16,700
S-1250	5	4	468	19,300
S-1500	6	4	524	21,900
S-1750	7	5	604	24,900
S-2000	8	5	660	27,500
S-4000	16	8	1156	48,700

* Without input/output.

Although up to three arrays can be packaged in one cabinet, the RADCAP configuration has two arrays per cabinet for symmetry. Figure 12 shows the equipment that was delivered to RADC. This includes a sequential control cabinet, an AP control cabinet, two AP memory cabinets for the four associative arrays, and a CIOU cabinet. The disk drive and line printer are mounted in separate cabinets. The keyboard/printer, the card reader, and the graphics display console can be mounted on table tops or pedestals. As mentioned earlier, the internal performance monitor is packaged within the CIOU cabinet. The external perfor-

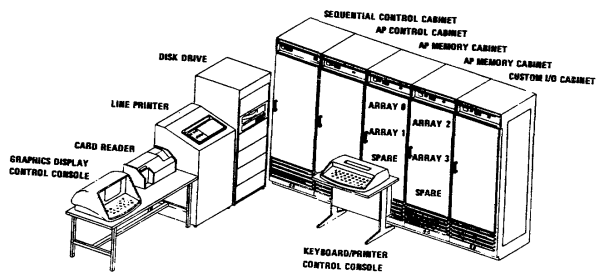

Figure 12—STARAN complex at RADC

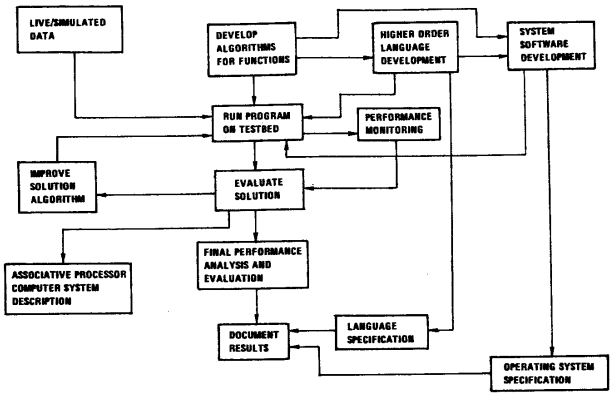

Figure 13—Flow of RADCAP research project

mance monitor, not shown in Figure 12, mounts on a table top.

Summary of system software

The system software available for STARAN/RADCAP is capable of operating STARAN in a stand-alone mode or, when integrated with the 645, in a STARAN/Multics configuration. The system software is based upon a disk operating system, which provides ready access to system programs, device independent I/O, and a file system. Operation of STARAN can be under direct control of the user at the control console or run in a batch mode with a control stream from an input device like the card reader.

The total assembly package for STARAN has a macro-language processor, an APPLE assembler, and a relocating linker. Programs are written in the APPLE and MAPPLE languages. Extensive string handling and substitution are implemented in the macro-preprocessor. APPLE is a symbolic language that includes mnemonics for parallel and associative operations. The linker combines separately assembled object modules by relocating code as necessary and resolving globally defined symbols.

Control of processing in STARAN is through interactive system routines. These routines are the interface between application program execution and the user. They allow the user to start and halt STARAN, to load programs and overlays, and to debug programs with trace, memory modification, and dump commands.

Diagnostic programs for STARAN hardware are disk resident. The programs can be called individually, in groups related to specific parts of the hardware, or as a total set for complete system testing. Fault detection and location are provided.

Additional software for the integrated STARAN/Multics operation is designed to handle the interface between the computers and the use of STARAN from Multics. For the interface, a special device driver module has been added to the STARAN disk operating system. This driver is similar to drivers used for peripherals. It has been specialized for Multics and can accommodate 16 open files simultaneously. A device interface module (DIM) has been added to Multics as the counterpart to the device driver. These two modules are basic parts of each machine's operating system and are transparent to the programmer.

STARAN can be operated from Multics by commands a user inputs at a terminal or from a file. File control procedures handle STARAN related keyboard inputs and provide the interface between the DIM and the MULTICS storage system. With these procedures, a user process executing in the 645 can call for execution of a STARAN program.

To facilitate the assembly of STARAN programs, a cross assembler is provided for time-shared use in Multics. This assembler accepts MAPPLE and APPLE as inputs.

Objectives and uses

The basic objective of the RADCAP facility is to explore the performance of a hybrid computer configuration (STARAN associative processor interfaced with a 645 sequential processor) on real-world, real-time problems. A specific goal is to determine the cost-effectiveness of associative/parallel processing in such an environment. Associative processing has been studied extensively in both theoretical and simulation studies, but no significant practical operating experience with them exists. Experimentation is necessary to provide "hard" data and fill in the presently existing void. Practical operating experience also is required so that a general-purpose associative processor configuration could be developed if results warrant it.

The RADCAP facility will be used in an experimental program to evaluate the internal performance of this hybrid computer configuration by means of hardware and/or software performance monitors to determine internal component utilization and system bottlenecks. Programming aspects of associative processing also will be investigated. Associative-processing programming is not well understood and represents radical departures from the traditional programming approach. The program loop is being replaced by hardware processing elements. This requires a whole new programming attitude. Programming languages suitable for associative processors probably will be quite different from present ones. This basic uncertainty must be explored and some practical operating experience gained. As a test problem, indicative of high data rate and real-time processing requirements, the data processing functions of an air surveillance system (AWACS) have been chosen. The primary functions to be investigated are tracking (both passive and active), display processing, and weapons control.

The scope of the research program can be described with the aid of Figure 13. The flow will begin with the development of associative-sequential algorithms for each of the AWACS data processing functions. As these algorithms are being developed, the application engineers will make known to a language and system software group those instruction level and system routine functions required to support the AWACS processing functions.

Based on this input, the language group will develop a language and implement this language on the RADCAP testbed. The system software activity will implement routines to support the command language. The applications program will then be run on the testbed using, where possible, nonsynthesized data as input. The machine activity will be monitored to gather statistics on utilization, identify system bottlenecks, and determine the efficiency with which the algorithms provide solution.

The data collected will then be analyzed to determine where cost-effective improvements can be made to software and/or hardware in order to improve the cost-effective performance of the system. These changes will be incorporated into the system via micro-program or software routines. If the change is to be a hardware design, that design will be made to the gate level so that performance and cost-effectiveness determination can be made.

When the solution to the problem is finally refined, it will be contrasted with known sequential solutions.

Initially each of the AWACS data processing functions will be treated separately. The final task will then be to develop a system executive and integrate all the functions to reflect the real world.

REFERENCES

1. Batcher, K. E., *STARAN Parallel Processor System Hardware,* GER-15996, Akron, Ohio, Goodyear Aerospace Corporation, 19 November 1973.
2. Davis, E. W., *STARAN Parallel Processor System Software,* GER-15997, Akron, Ohio, Goodyear Aerospace Corporation, Akron, Ohio, 19 November 1973.
3. Rudolph, J. A., "A Production Implementation of an Associative Array Processor—STARAN," *1972 Fall Joint Computer Conference Proceedings,* December 1972, pp. 229–241.
4. Slade, A. E. and H. O. McMahon, "The Cryatron Catalog Memory System," *1957 Fall Joint Computer Conference Proceedings,* Vol. 10, pp. 115–120.
5. Fulmer, L. C. and W. C. Meilander, "A Modular Plated Wire Associative Processor," *Proceedings of IEEE Computer Group Conference,* June 1970.
6. Shore, J. E., "Second Thoughts on Parallel Processing," *NRL Report 7364,* December 30, 1971, page 7.
7. Kressler, R. R., C. E. Peet, Jr., and F. B. Frazee, "Development of an LSI Associative Processor," *Air Force Avionics Laboratory Technical Report 70-142,* August 1970.
8. Boyle, J. T. and C. A. Neugebauer, "Approaches to Custom Large Scale Integration," *Air Force Avionics Laboratory Technical Report 73-66,* March 1973.
9. Batcher, K. E., "Flexible Parallel Processing and STARAN," *1972 WESCON Technical Papers,* Session 1.