

Huppenthal

Reference 1

We believe the key to the 10's longevity is its basically simple, clean structure with adequately large (one Mbyte) address space that allows users to get work done. In this way, it has evolved easily with use and with technology. An equally significant factor in its success is a single operating system environment enabling user program sharing among all machines. The machine has thus attracted users who have built significant languages and applications in a variety of environments. These user-developers are thus the dominant system architects-implementors.

In retrospect, the machine turned out to be larger and further from a minicomputer than we expected. As such it could easily have died or destroyed the tiny DEC organization that started it. We hope that this paper has provided insight into the interactions of its development.

Acknowledgments. Dan Siewiorek deserves our greatest thanks for helping with a complete editing of the text. The referees and editors have been especially helpful. The important program contributions by users are too numerous for us to give by name but here are most of them: APL, Basic, BLISS, DDT, LISP, Pascal, Simula, SOS, TECO, and Tenex. Likewise, there have been so many contributions to the 10's architecture and implementations within DEC and throughout the user community that we dare not give what would be a partial list.

Received April 1977; revised September 1977

References

1. Bell, G., Cady, R., McFarland, H., Delagi, B., O'Laughlin, J., and Noonan, R. A new architecture for minicomputers—the DEC PDP-11. Proc. AFIPS 1970 SJCC, Vol. 36, AFIPS Press, Montvale, N.J., pp. 657-675.
2. Bell, G., and Freeman, P. Cai—A computer architecture for AI research AFIPS Conf. Proc. Vol. 38 (Spring, 1971), 779-790.
3. Bell, G., and Newell, A. *Computer Structures: Readings and Examples*. McGraw-Hill, New York, 1971.
4. Bobrow, D.G., Burchfiel, J.D., Murphy, D. L., and Tomlinson, R.S. TENEX, A Paged Time Sharing System for the PDP-10. *Comm. ACM* 15, 3 (March 1972), 135-143.
5. Bullman, D.M. Editor, stack computers issue. *Computer* 10, 5 (May 1977), 14-52.
6. Clark, W.A. The Lincoln TX-2 computer. Proc. WJCC 1957, Vol. 11, pp. 143-171.
7. Lunde, A. Empirical evaluation of some features of Instruction Set Processor architecture. *Comm. ACM* 20, 3 (March 1977), 143-152.
8. Mitchell, J.L., and Olsen, K.H. TX-0, a transistor computer. Proc. EJCC 1956, Vol. 10, pp. 93-100.
9. McCarthy, J. *Time Sharing Computer Systems, Management and the Computer of the Future* M. Greenberger, Ed., M.I.T. Press, Cambridge, Mass., 1962, pp. 221-236.
10. Murphy, D.L. Storage organization and management in TENEX. Proc. AFIPS 1972 FJCC, Vol. 41, Pt. I, AFIPS Press, Montvale, N.J., pp. 23-32.
11. Olsen, K.H. Transistor circuitry in the Lincoln TX-2. Proc. WJCC 1957, Vol. 11, pp. 167-171.
12. Roberts, L.G. Ed. Section on Resource Sharing Computer Networks. AFIPS 1970 SJCC, Vol. 36, AFIPS Press, Montvale, N.J., pp. 543-598.
13. Wulf, W., and Bell, G. C.mmp—A multi-mini-processor. Proc. AFIPS 1972 FJCC, Vol. 41, AFIPS Press, Montvale, N.J., pp. 765-777.
14. Wulf, W., Russell, D., and Habermann, A.N. BLISS: A language for systems programming. *Comm. ACM* 14, 12 (Dec. 1971), 780-790.

63

Computer
Systems

G. Bell, S. H. Fuller, and
D. Siewiorek, Editors

The CRAY-1 Computer System

Richard M. Russell
Cray Research, Inc.

This paper describes the CRAY-1, discusses the evolution of its architecture, and gives an account of some of the problems that were overcome during its manufacture.

The CRAY-1 is the only computer to have been built to date that satisfies ERDA's Class VI requirement (a computer capable of processing from 20 to 60 million floating point operations per second) [1].

The CRAY-1's Fortran compiler (CRT) is designed to give the scientific user immediate access to the benefits of the CRAY-1's vector processing architecture. An optimizing compiler, CRT, "vectorizes" innermost DO loops. Compatible with the ANSI 1966 Fortran Standard and with many commonly supported Fortran extensions, CRT does not require any source program modifications or the use of additional nonstandard Fortran statements to achieve vectorization. Thus the user's investment of hundreds of man months of effort to develop Fortran programs for other contemporary computers is protected.

Key Words and Phrases: architecture, computer systems

CR Categories: 1.2, 6.2, 6.3

Introduction

Vector processors are not yet commonplace machines in the larger-scale computer market. At the time of this writing we know of only 12 non-CRAY-1 vector processor installations worldwide. Of these 12, the most powerful processor is the ILLIAC IV (1 installation), the most populous is the Texas Instruments Advanced Scientific Computer (7 installations) and the most publicized is Control Data's STAR 100

Copyright © 1977, Association for Computing Machinery, Inc. General permission to republish, but not for profit, all or part of this material is granted provided that ACM's copyright notice is given and that reference is made to the publication, to its date of issue, and to the fact that reprinting privileges were granted by permission of the Association for Computing Machinery.

Author's address: Cray Research Inc., Suite 213, 7850 Metro Parkway, Minneapolis, MN 55420.

Communications
of
the ACM

January 1978
Volume 21
Number 1

(4 installations). In its report on the CRAY-1, Auerbach Computer Technology Reports published a comparison of the CRAY-1, the ASC, and the STAR 100 [2]. The CRAY-1 is shown to be a more powerful computer than any of its main competitors and is estimated to be the equivalent of five IBM 370/195s.

Independent benchmark studies have shown the CRAY-1 fully capable of supporting computational rates of 138 million floating-point operations per second (MFLOPS) for sustained periods and even higher rates of 250 MFLOPS in short bursts [3, 4]. Such comparatively high performance results from the CRAY-1 internal architecture, which is designed to accommodate the computational needs of carrying out many calculations in discrete steps, with each step producing interim results used in subsequent steps. Through a technique called "chaining," the CRAY-1 vector functional units, in combination with scalar and vector registers, generate interim results and use them again immediately without additional memory references, which slow down the computational process in other contemporary computer systems.

Other features enhancing the CRAY-1's computational capabilities are: its small size, which reduces distances electrical signals must travel within the computer's framework and allows a 12.5 nanosecond clock period (the CRAY-1 is the world's fastest scalar processor); a one million word semiconductor memory equipped with error detection and correction logic (SECDED); its 64-bit word size; and its optimizing Fortran compiler.

Architecture


The CRAY-1 has been called "the world's most expensive love-seat" [5]. Certainly, most people's first reaction to the CRAY-1 is that it is so small. But in computer design it is a truism that smaller means faster. The greater the separation of components, the longer the time taken for a signal to pass between them. A cylindrical shape was chosen for the CRAY-1 in order to keep wiring distances small.

Figure 1 shows the physical dimensions of the machine. The mainframe is composed of 12 wedge-like columns arranged in a 270° arc. This leaves room for a reasonably trim individual to gain access to the interior of the machine. Note that the love-seat disguises the power supplies and some plumbing for the Freon cooling system. The photographs (Figure 2 and 3) show the interior of a working CRAY-1 and an exterior view of a column with one module in place. Figure 4 is a photograph of the interior of a single module.

An Analysis of the Architecture

Table I details important characteristics of the CRAY-1 Computer System. The CRAY-1 is equipped with 12 i/o channels, 16 memory banks, 12 functional

Fig. 1. Physical organization of mainframe.



- Dimensions
 - Base-103½ inches diameter by 19 inches high
 - Columns-56½ inches diameter by 77 inches high including height of base
- 24 chassis
- 1662 modules; 113 module types
- Each module contains up to 288 IC packages per module
- Power consumption approximately 115 kw input for maximum memory size
- Freon cooled with Freon/water heat exchange
- Three memory options
- Weight 10,500 lbs (maximum memory size)
- Three basic chip types
 - 5/4 NAND gates
 - Memory chips
 - Register chips

units, and more than 4k bytes of register storage. Access to memory is shared by the i/o channels and high-speed registers. The most striking features of the CRAY-1 are: only four chip types, main memory speed, cooling system, and computation section.

Four Chip Types

Only four chip types are used to build the CRAY-1. These are 16 × 4 bit bipolar register chips (6 nanosecond cycle time), 1024 × 1 bit bipolar memory chips (50 nanosecond cycle time), and bipolar logic chips with subnanosecond propagation times. The logic chips are all simple low- or high-speed gates with both a 5 wide and a 4 wide gate (5/4 NAND). Emitter-coupled logic circuit (ECL) technology is used throughout the CRAY-1.

The printed circuit board used in the CRAY-1 is a 5-layer board with the two outer surfaces used for signal runs and the three inner layers for -5.2V, -2.0V, and ground power supplies. The boards are six inches wide, 8 inches long, and fit into the chassis as shown in Figure 3.

All integrated circuit devices used in the CRAY-1 are packaged in 16-pin hermetically sealed flat packs supplied by both Fairchild and Motorola. This type of package was chosen for its reliability and compactness. Compactness is of special importance; as many as 288 packages may be added to a board to fabricate a module (there are 113 module types), and as many as 72 modules may be inserted into a 28-inch-high chassis.

Fig. 2. The CRAY-1 Computer.





Fig. 3. CRAY-1 modules in place.



Such component densities inevitably lead to a mammoth cooling problem (to be described).

Main Memory Speed

CRAY-1 memory is organized in 16 banks, 72 modules per bank. Each module contributes 1 bit to a 64-bit word. The other 8 bits are used to store an 8-bit check byte required for single-bit error correction, double-bit error detection (SEDED). Data words are stored in 1-bank increments throughout memory. This organization allows 16-way interleaving of memory accesses and prevents bank conflicts except in the case

Fig. 4. A single module.




Table I. CRAY-1 CPU characteristics summary

Computation Section

- Scalar and vector processing modes
- 12.5 nanosecond clock period operation
- 64-bit word size
- Integer and floating-point arithmetic
- Twelve fully segmented functional units
- Eight 24-bit address (*A*) registers
- Sixty-four 24-bit intermediate address (*B*) registers
- Eight 64-bit scalar (*S*) registers
- Sixty-four 64-bit intermediate scalar (*T*) registers
- Eight 64-element vector (*V*) registers (64-bits per element)
- Vector length and vector mask registers
- One 64-bit real time clock (*RT*) register
- Four instruction buffers of sixty-four 16-bit parcels each
- 128 basic instructions
- Prioritized interrupt control

Memory Section

- 1,048,576 64-bit words (plus 8 check bits per word)
- 16 independent banks of 65,536 words each
- 4 clock period bank cycle time
- 1 word per clock period transfer rate for *B*, *T*, and *V* registers
- 1 word per 2 clock periods transfer rate for *A* and *S* registers
- 4 words per clock period transfer rate to instruction buffers (up to 16 instructions per clock period)

i/o Section


- 24 i/o channels organized into four 6-channel groups
- Each channel group contains either 6 input or 6 output channels
- Each channel group served by memory every 4 clock periods
- Channel priority within each channel group
- 16 data bits, 3 control bits per channel, and 4 parity bits
- Maximum channel rate of one 64-bit word every 100 nanoseconds
- Maximum data streaming rate of 500,000 64-bit words/second
- Channel error detection

of memory accesses that step through memory with either an 8 or 16-word increment.

Cooling System

The CRAY-1 generates about four times as much heat per cubic inch as the 7600. To cool the CRAY-1 a new cooling technology was developed, also based on Freon, but employing available metal conductors in a new way. Within each chassis vertical aluminum/stainless steel cooling bars line each column wall. The

Fig. 5. Block diagram of registers.



Freon refrigerant is passed through a stainless steel tube within the aluminum casing. When modules are in place, heat is dissipated through the inner copper heat transfer plate in the module to the column walls and thence into the cooling bars. The modules are mated with the cold bar by using stainless steel pins to pinch the copper plate against the aluminum outer casing of the bar.

To assure component reliability, the cooling system

was designed to provide a maximum case temperature of 130°F (54°C). To meet this goal, the following temperature differentials are observed:

Temperature at center of module	130°F (54°C)
Temperature at edge of module	118°F (48°C)
Cold plate temperature at wedge	78°F (25°C)
Cold bar temperature	70°F (21°C)
Refrigerant tube temperature	70°F (21°C)

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.