

CRAY-1 Computer Technology

JAMES S. KOLODZEY, MEMBER, IEEE

Abstract—Hardware and packaging technology which provide the high performance of the CRAY-1 computer are reviewed. A brief overview of the computer is given, followed by a description of the computer circuits, packaging, power distribution, and cooling system.

I. INTRODUCTION

SINCE ITS introduction in 1976, the CRAY-1 has developed a reputation as a fast and reliable scientific processor. The CRAY-1S, announced in 1979, offers enhanced input/output (I/O) capability and an increase of the maximum memory size from one million to four million words. A photo of the CRAY-1S is shown in Fig. 1, where the large section at the left contains the central processing unit (CPU) and memory, and the smaller section at the right contains the I/O processor. Fig. 2 shows the CPU chassis layout and circuit module locations. Data and control signal paths are given in Fig. 3. Applications for the CRAY-1 include large-scale calculations of the type required in weather forecasting, petroleum and earthquake seismology, structural analysis, nuclear engineering, and particle physics.

The CRAY-1 characteristic of combined high performance and reliability results from simple hardware design. One chip type is used predominantly in the CPU, and one memory chip type is used in the memory banks. Each arithmetic operation has a dedicated hardware functional unit, and the logic and circuit design is governed by straightforward ground rules. The power supplies are simple rectifier/filter types and the cooling system is a standard refrigeration unit. Whenever possible problems were solved with existing technology. This paper reviews the CRAY-1 hardware in detail.

II. SYSTEM OVERVIEW

Table I is a listing of some CRAY-1S performance characteristics. The logic gates in the CPU are mostly emitter-coupled logic (ECL) dual NAND integrated circuits (IC's). Gate counts of the hardware functional units are given in Table II. Additional gates are used for control and storage functions which tie the hardware functions together. A full memory CRAY-1S contains four million words of high speed ECL random access memory (RAM). The $4K \times 1$ bit random access memory (RAM) chips are arranged to give a word length of 64 bits plus 8 bits for single error correction, double error detection (SECDED). The SECDED is an implementation of a Hamming code. A total of 73 728 RAM chips are used.

The CRAY-1 is a vector processor with the ability to operate iteratively on strings of up to 64 (vector length) operands




Fig. 1. CRAY-1S mainframe.

or operand pairs. By contrast, scalar processors perform one iteration on one or a pair of operands. Each operand has a word length of 64 bits, which for floating point operations comprises a 49-bit signed coefficient and a 15-bit exponent. The CRAY-1 has the ability to operate efficiently even with short vector lengths [1], and it is also a fast scalar processor.

Machine performance is expressed in millions of floating point operations per second (megaflops) because a single vector instruction is equivalent to a loop of several scalar instructions. The CRAY-1 has been shown capable of a sustained rate of 138 Mflops and to achieve 250 Mflops in short bursts [2]. System reliability is excellent, with an availability greater than 98 percent. Mean time between interruption (MTBI) is more than 100 h, and mean time to repair (MTTR) is about 1 h [3].

III. CIRCUIT COMPONENTS

One chip type comprises about 95 percent of the IC's in the CPU. The chip is a negative logic, five and four input dual NAND gate (5/4 gate) with complementary outputs. A logic diagram is shown in Fig. 4. The 5/4 gate is an ECL circuit with 750 ps propagation delay and 60 mW per gate power dissipation (120 mW per package). The 0.036 in square silicon die is packaged in a hermetically sealed 16-pin ceramic flatpack. Package dimensions are $3/8 \times 1/4 \times 1/12$ in with leads on 50 mil centers as shown along with the smaller resistor package (described below) in Fig. 5. The ceramic flatpack is used for reliability, small size, and speed.

Other chips are a dual *D* flip flop, a 16×4 bit register with a 6 ns access time, and the $4K \times 1$ bit RAM with 25 ns access time packaged in an 18-pin ceramic flatpack.

In addition to the IC's the only other circuit component on the printed circuit (PC) board is the transmission line termina-

Manuscript received October 16, 1980; revised January 29, 1981.

The author is with Cray Laboratories, 3375 Mitchell Lane, Boulder, CO 80301.

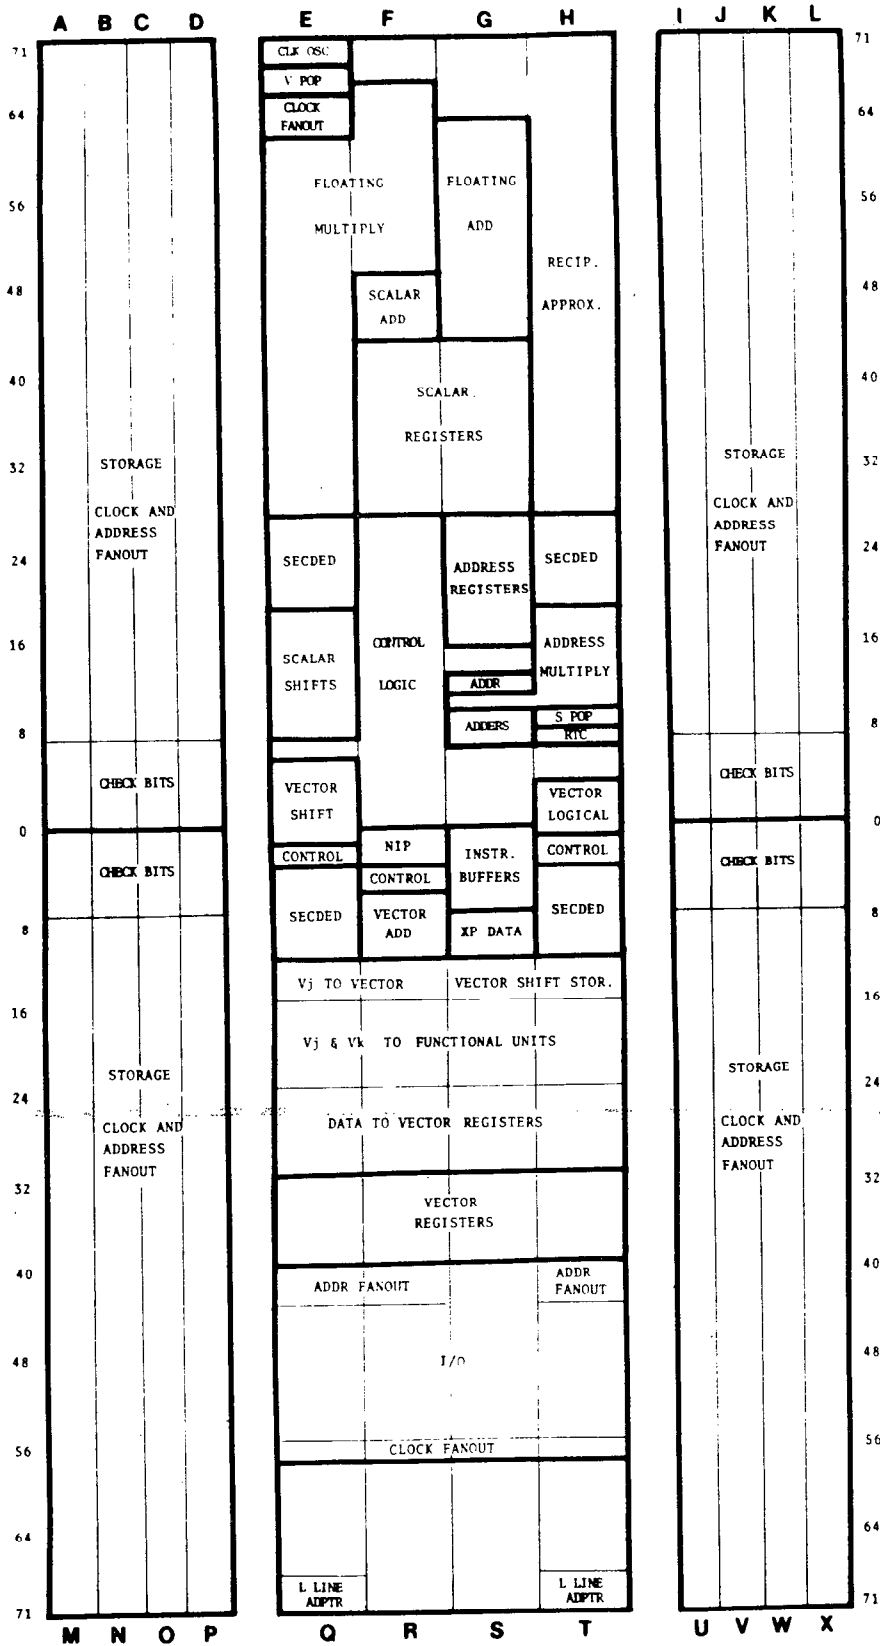


Fig. 2. General chassis layout.

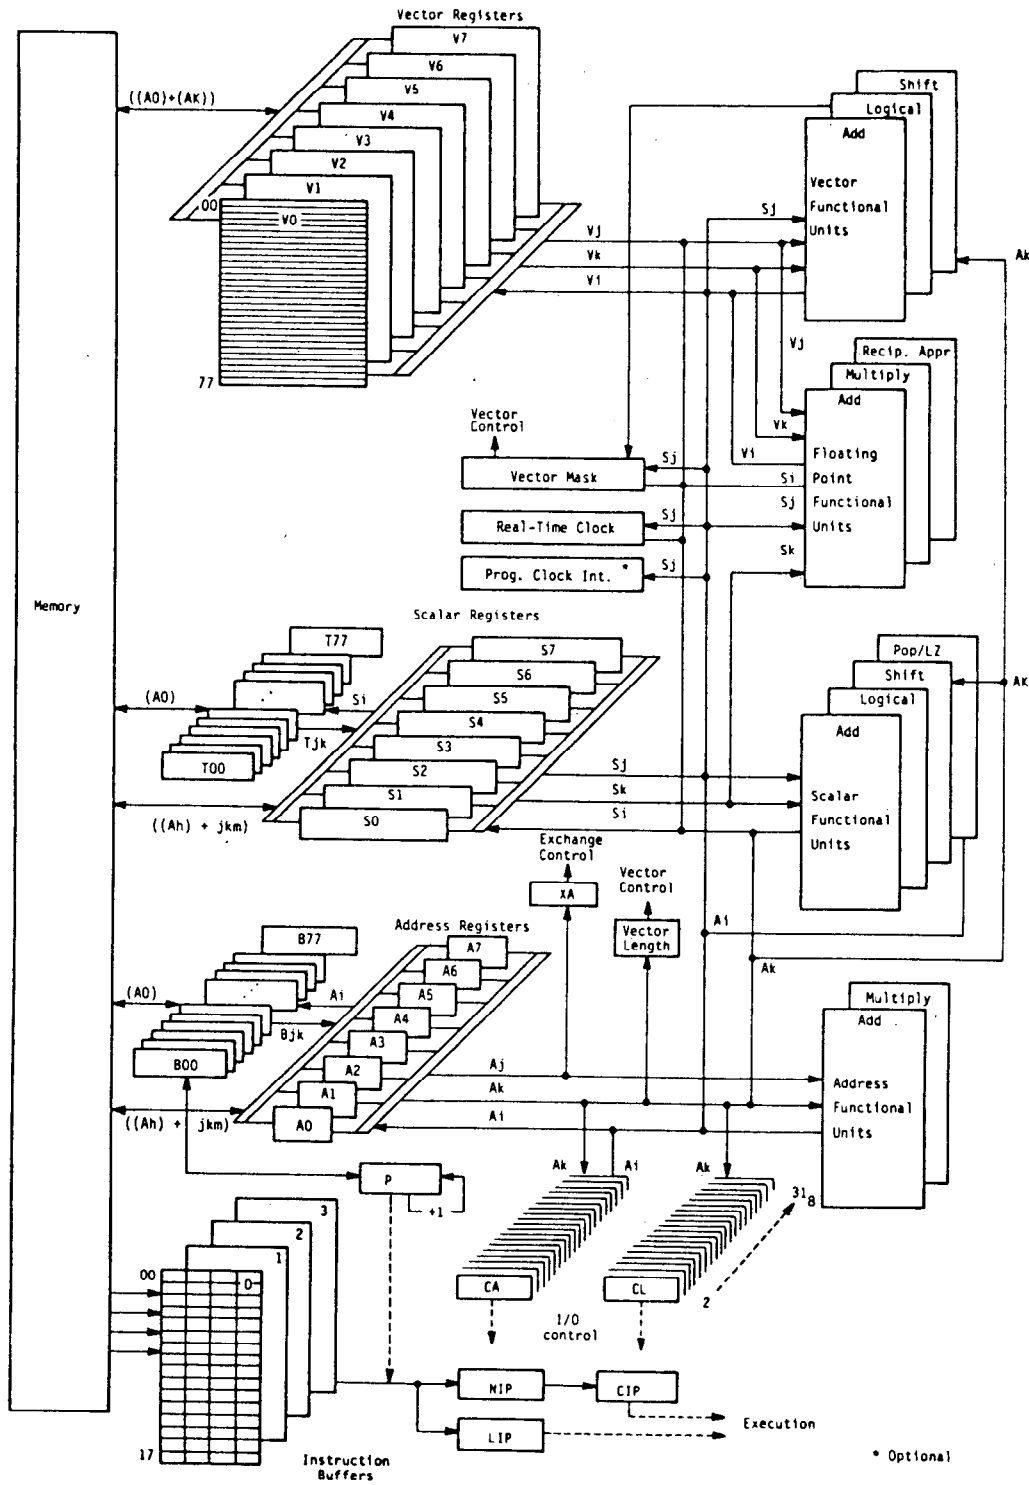


Fig. 3. Computation section.

TABLE I
CRAY-1S SPECIFICATIONS

| | |
|---------------|------------------|
| CPU: | 230 000 gates |
| Memory: | 4 million words |
| Clock: | 12.5 ns |
| Power: | 130 kW |
| Operations: | 138 Mflops |
| | 250 Mflops burst |
| Availability: | 98 percent |
| MTBI: | 100 h |

TABLE II
CRAY-1S GATE COUNTS

| Functional Unit | Gate Count | Percentage of Total |
|--------------------------|------------|---------------------|
| Address adder | 1952 | 2.59 |
| Address multiply | 4009 | 5.33 |
| Scalar add | 2968 | 3.94 |
| Scalar single shift | 1452 | 1.93 |
| Scalar double shift | 2976 | 3.95 |
| Constant to Si | 482 | 0.64 |
| Pop and zero count to Ai | 403 | 0.54 |
| Vector integer add | 2216 | 2.94 |
| Vector logical | 1984 | 2.64 |
| Vector shift | 3460 | 4.60 |
| Vector pop count | 490 | 0.65 |
| Floating add | 8247 | 11.0 |
| Floating multiply | 23116 | 30.7 |
| Reciprocal | 21504 | 28.6 |
| Total | 75259 | 100.05 |




Fig. 4. 5/4 gate schematic.




Fig. 5. 5/4 gate and resistor package.

tion resistor. Two 60- Ω tantalum nitride thin film resistors are packaged in a ceramic *T* package with a common termination voltage lead. The 1300 W/in² power handling capability of the thin film material allows for a very small size resistor. Peak power dissipation is approximately 20 mW per resistor,

resulting in about 1.5 W/in² for the resistor film. A total of 495 934 resistor packages are used in the CRAY-1.

IV. CIRCUIT INTERCONNECTION

The field replaceable units in the CRAY-1 are modules of the type shown in Fig. 6. Visible here is a 6 \times 8 in multilayer printed circuit board which can accommodate up to 144 IC packages in a 12 \times 12 array. The PC boards are five-layer structures with the following layer assignments: components and signal traces; ground (V_{CC}); $-2V(V_{TT})$; $-5.2V(V_{EE})$; and a second-signal layer. The board material is G10 glass-epoxy with 1 oz copper conductor layers. Via and component holes are plated through with a 0.022-in inside diameter.

With signal rise times (10 percent-90 percent) of 750 ps, open line stub lengths must be less than 0.5 in to hold reflections under 35 percent overshoot and under 12 percent undershoot. Larger reflections can saturate gate inputs and reduce noise immunity. Longer signal runs must use transmission lines terminated in the line characteristic impedance. The CRAY-1 boards have 7-mil wide lines and 7-mil spaces. The 7-mil height above a ground (or power) plane results in a 60- Ω microstrip line with a delay of 0.15 ns/in. Gate loads appear as open stubs on the line as shown in Fig. 7. Fan out is limited to four on board or three off board to reduce ac loading problems.

Signal propagation between gates is governed by strict timing rules as shown in Fig. 8. The 12.5 ns clock is divided into eight "gate times" of about 1.5 ns each. Roughly half the gate time is due to circuit propagation delay, and half is due to board-foil delay. Unused gates can be dropped from the path by adding 3 in of foil conductor.

A CRAY-1 module consists of two PC boards sandwiching a 0.08-in thick copper cooling plate that is also the ground bus. Signal communication between boards is performed differentially over 120- Ω twisted pair in the backplane as shown schematically in Fig. 9. The 120- Ω twisted pair is matched to two 60- Ω board traces in series. This differential method permits a low cross talk communication path and provides both the signal and its complement for use on the receiving module. An inverter gate (and gate delay) is therefore saved if the signal complement is needed. The twisted pair attaches to the module by a 96-pin pair connector with pins on 0.05 in centers. The board side of the connector pair is visible in Fig. 6 at the top.

Interconnection ground rules restrict twisted pair lengths to multiples of gate times, or multiples of 1 ft to a maximum of 4 ft. These circuit rules prevent timing problems or race conditions from happening anywhere in the machine. The CRAY-1 backplane contains 67 mi of twisted pair wire. The first machine was completed with no wiring errors.

V. SYSTEM POWER

Voltage requirements for the PC board are $-5.2V(V_{EE})$ for the IC's and $-2V(V_{TT})$ for the termination resistors. No bypass capacitors are used on the PC board. A benefit of using the 5/4 gate IC is that it provides a balanced load to the supply. When one output turns on, another turns off, and the power supply loading is purely resistive and constant. Any ripple which occurs during the transitions is filtered out by




Fig. 6. CRAY-1 module.




Fig. 9. Intermodule communication.

the 16-nF capacitor formed by the power and ground plane in the PC board.

A CPU module (two boards) can dissipate up to 36 W (7 A) from V_{EE} and about 1.2 W (0.6 A) from V_{TT} . There are 576 CPU modules. Taking an average of 25 W per module, the CPU dissipates a total of 14.4 kW. Each memory module contains 64 of the 4K × 1 bit ECL RAM chips which dissipate about 1 W each, and some interface logic chips. The power dissipation of the memory module is about 70 W. A full four million word CRAY-1S contains 1152 memory modules with a total memory power dissipation of 81 kW. Total power dissipation for the computer is 95 kW. Approximately 130 kW is supplied to the entire machine including power supply losses.

The computer modules receive V_{EE} and V_{TT} power from power supplies located under the seats around the periphery of the computer. These power supplies are simple linear rectifier-filter types which receive 400 Hz ac voltage from 36 variable transformers on a power distribution unit (PDU). The PDU receives 208 Vac at 400 Hz from a motor generator (MG) set which converts 480 V at 60 Hz to 208 V at 400 Hz. Each installation has two MG sets with one available for backup.

VI. SYSTEM COOLING

The CRAY-1S cooling system is designed to limit the IC die temperature to a maximum of 65°C. This provides a reliability margin from the 150°C absolute maximum IC junction temperature. The IC package case is maintained at 54°C. Heat generated in the silicon die flows through the IC package to the PC board ground plane and then to the 0.08-in thick copper cold plate. The cold plate conducts the board heat to its edges, which are held to 25°C by contact with a cast aluminum cold bar. The aluminum cold bars form the twelve vertical columns in the computer mainframe into which the modules slide horizontally on 0.4-in spacings. A refrigerant, Freon 22, flows through stainless steel tubes embedded in the cold bars. The development of the composite aluminum/stainless steel cold bars represents a solution to one of the more difficult design problems of the CRAY-1. Cast aluminum is actually porous and oil mixed in with the Freon can cause reliability problems if it leaks onto the modules. A method to bond stainless steel tubing into cast aluminum had to be invented in order to make the cold bars practical.

The refrigerant is maintained at 18.5°C by an evaporative refrigeration system. Freon 22, which boils -41°C (at atmospheric pressure), absorbs heat from the cold bar and changes to the gas phase. It passes through a compressor and condenses

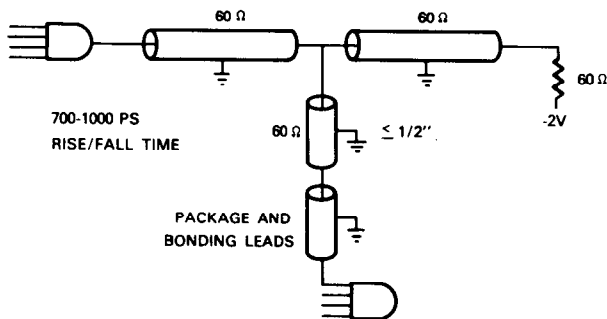


Fig. 7. Printed circuit transmission lines.




Fig. 8. CRAY-1 gate timing.

back to a liquid by releasing heat to a cold water supply which flows at 40 gal/min. The maximum heat load of this system is 580 000 B/h.

VII. FUTURE TECHNOLOGY

Key hardware problems in designing high-speed computers are heat removal, circuit interconnection, IC packaging, circuit density, and device speed. Improvements in these areas can directly improve computer performance.

Many factors contribute to the cooling problem. A given circuit technology is power-delay product limited. Short gate propagation delays require a high gate power dissipation. Large numbers of gates per chip as in LSI will cause total chip power to be well over 1 W, and the requirement for dense packaging for short wire delays will cause a heat density problem. Furthermore, cooler circuits run more reliably. All this creates demand for a high performance cooling system.

Until the day when a whole mainframe is on a chip, interchip wire propagation delays will slow down system speed. This delay is minimized by using low dielectric PC board materials such as Teflon ($\epsilon_r = 2.4$) and polyimide ($\epsilon_r = 3.5$) rather than glass-epoxy ($\epsilon_r = 4.5$). Using surface microstrip lines rather than buried stripline also decreases wire delay. Fine signal line widths and spacings help shrink chip-chip spacings, especially with large numbers of leads per package. Present CRAY-1 boards use 7 mil lines and spaces. Four to five mil lines may be the limit for subtractively etched PC boards, so new techniques will be needed.

Smaller IC packages help both mechanically with reduced PC board real-estate and electrically with reduced lead inductance and capacitance. Table III lists some properties of packaging materials. Improved designs such as ceramic chip carriers permit a higher board density and a more uniform wire length from external lead to IC die. A rough rule of thumb for wire inductance is 12 nH/in. With subnanosecond signal edge speeds, bond wire lengths of more than a few hundred mils can cause severe delay and waveform degradation. A package can add 1 pF or more to the gate input capacitance. This can add to signal delay by a factor of 30 ps/pF.

Improved circuit technology gives a direct benefit to computer speed. Table IV gives some rough numbers of circuit performance for comparison purposes. Silicon ECL may reach a delay limit in the low hundreds of picoseconds, and a power-delay limit in the low picojoule range. Significant improvement is gained by switching to high mobility materials such as gallium arsenide or indium phosphide. Technical problems need to be solved, however. Processing of these materials is still in the early stages. Gates have been fabricated using GaAs transistors, but problems exist with device threshold voltage uniformity and limited current drive due to input saturation as compared to silicon bipolar or MOS devices. An insulated gate structure would alleviate the current drive

TABLE III
SOME MATERIAL PARAMETERS

| | Signal Line Width | Dielectric Constant | Thermal Conductivity |
|-------------|-------------------|---------------------|---|
| Glass-Epoxy | 0.004 in | 4.5 | $0.0024 \frac{\text{W}}{\text{CM}^\circ\text{C}}$ |
| Alumina | 0.003 in | 8.9 | $0.28 \frac{\text{W}}{\text{CM}^\circ\text{C}}$ |
| Beryllia | 0.003 in | 6.8 | $1.68 \frac{\text{W}}{\text{CM}^\circ\text{C}}$ |
| Silicon | 2 μm | 11.7 | $1.50 \frac{\text{W}}{\text{CM}^\circ\text{C}}$ |

TABLE IV
CIRCUIT TECHNOLOGY GATE PROPAGATION DELAYS AND POWER DELAY PRODUCTS

| | T_{pd} | $P_d \times T_{pd}$ |
|---------------------|----------|---------------------|
| Silicon | 400 ps | 4.5 pJ |
| Gallium arsenide | 80 ps | 0.1 pJ |
| Josephson junctions | 15 ps | 1.0 fJ |

limitation and permit a high performance memory circuit. Further speed increase from these materials occurs by operating at liquid nitrogen temperatures (77 K) [4]. At low temperatures, device mobility increases and delays of 20 ps and power-delays of 10 fJ may be possible [5]. The lowest power technology would use Josephson Junctions (JJ's), but the necessity of liquid helium (4°K) cooling may make JJ's less attractive than cooled GaAs.

VIII. CONCLUSION

This paper has provided a review of the CRAY-1 computer hardware. It has been shown that high performance components, clever packaging, and systems design have made the CRAY-1 system practical [6].

ACKNOWLEDGMENT

The author wishes to thank all the Cray personnel whose efforts and contributions made the CRAY-1, and therefore this paper, possible.

REFERENCES

- [1] P. M. Johnson, "An introduction to vector processing," *Comput. Design*, pp. 89-97, Feb. 1978.
- [2] R. M. Russell, "The CRAY-1 Computer System," *Commun. Assoc. Comput. Machinery*, vol. 21, no. 1, pp. 63-72, Jan. 1978.
- [3] Cray Research Inc. field engineering statistics.
- [4] M. S. Shur and L. F. Eastman, "Ballistic and near ballistic transport in GaAs," *IEEE Trans. Electron Devices*, vol. ED-1, no. 8, pp. 147-148, Aug. 1980.
- [5] W. Twaddell, "IC's and semiconductors," *EDN*, pp. 37-50, Dec. 15, 1979.
- [6] *CRAY-1S Series Hardware Reference Manual*, CRAY-1 Computer Systems, HR-0808, Cray Research Inc., 1980.