INTRODUCTION

The Mark III^{fp} Hypercube and its operating system software represent the culmination of the efforts of the JPL/Caltech Campus collaboration started in 1983 with the design and construction of the Mark II 8086- based machine.

The Mark III, a Motorola 68020/68881-based machine, was first made operational in a prototype version in May 1985. JPL is currently adding a high-performance Weitek-based floating-point accelerator and will construct a single, 128-node, gigaflop machine during this year.

Mark III^{fp}

I. Hardware

The Mark III^{fp} is a new generation of the hypercube concurrent processing system developed at JPL/Caltech with peak performance of 5 MIPS, 14 MFLOPS per node, and a peak communication rate of 6 Mbytes per second. Each node board consists of two Motorola MC68020 processors, an MC 68881/68882 scalar floating-point coprocessor, and the Weitek WTL8000 processor chip set. One of the 68020s serves as the application and computational processor, while the other is dedicated for communication. The Weitek processor serves as a high-speed floating-point engine peripheral to the 68020 application processor. The three processors are interconnected through a common system bus and a shared 4 Mbytes of dynamic memory. Each processor has its own fast local memory, which increases parallelism, minimizes main memory referencing, reduces memory contentions, and improves system performance (See Figures 1, 1a & 1b).

The application processor has 128 Kbytes of local memory, the floating-point coprocessor, and a memory management unit for memory protection and multitasking applications. The communication processor contains special circuitry required to perform internal communications through one or more of the eight FIFO communications channels.

The shared memory is partitioned into four interleaved-address dynamic modules supporting all combinations of byte-, word- or long-word accesses, both aligned and misaligned. A 32-bit parallel error detection and correction circuit is available for single-bit detection/correction and double-bit detection.

The floating-point daughter board (FPDB), which uses the Weitek WTL8000 processor chip set, can provide up to 100X speedup for floating-point calculations compared to the Mark III node board's Motorola MC68020 processor with its 68881 floating-point coprocessor. The hardware/software architecture of the FPDB is such that the Weitek processor acts as the master or as a slave (subroutine engine) to the MC68020 processor. On-board instruction (128-Kbytes) and data (64-Kbytes) caches and bus arbitration circuitry allow the FPDB to access the node board's full 4-Mbyte memory space. A rich set of special features provides a high degree of control over cache operation, bus arbitration priority, and other functions, including 64 Kbytes of direct-mapped data RAM, allowing hand-coded applications to achieve maximum performance. (See Figure 2).

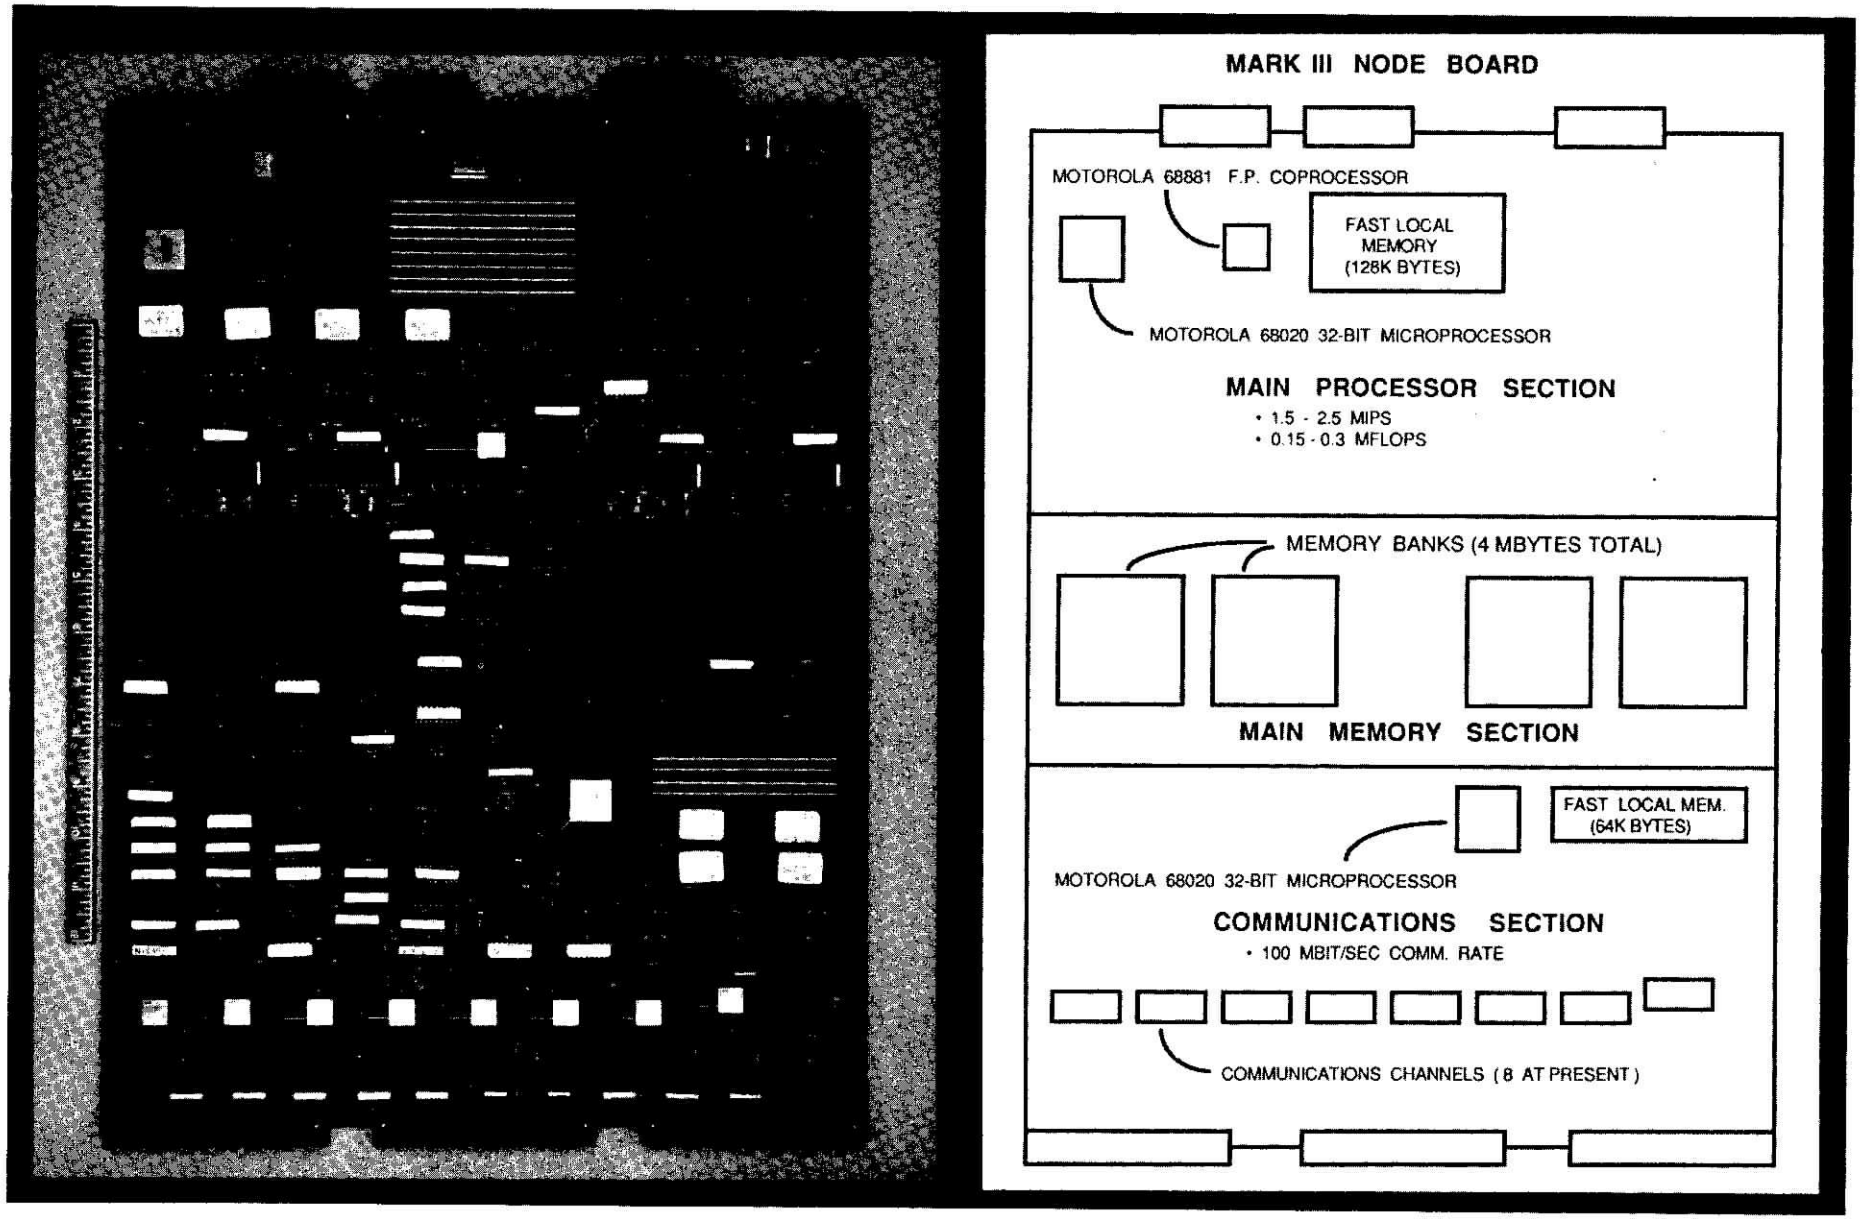

FIGURE 1 Photograph & Module Identification of the Mark III Hypercube Node

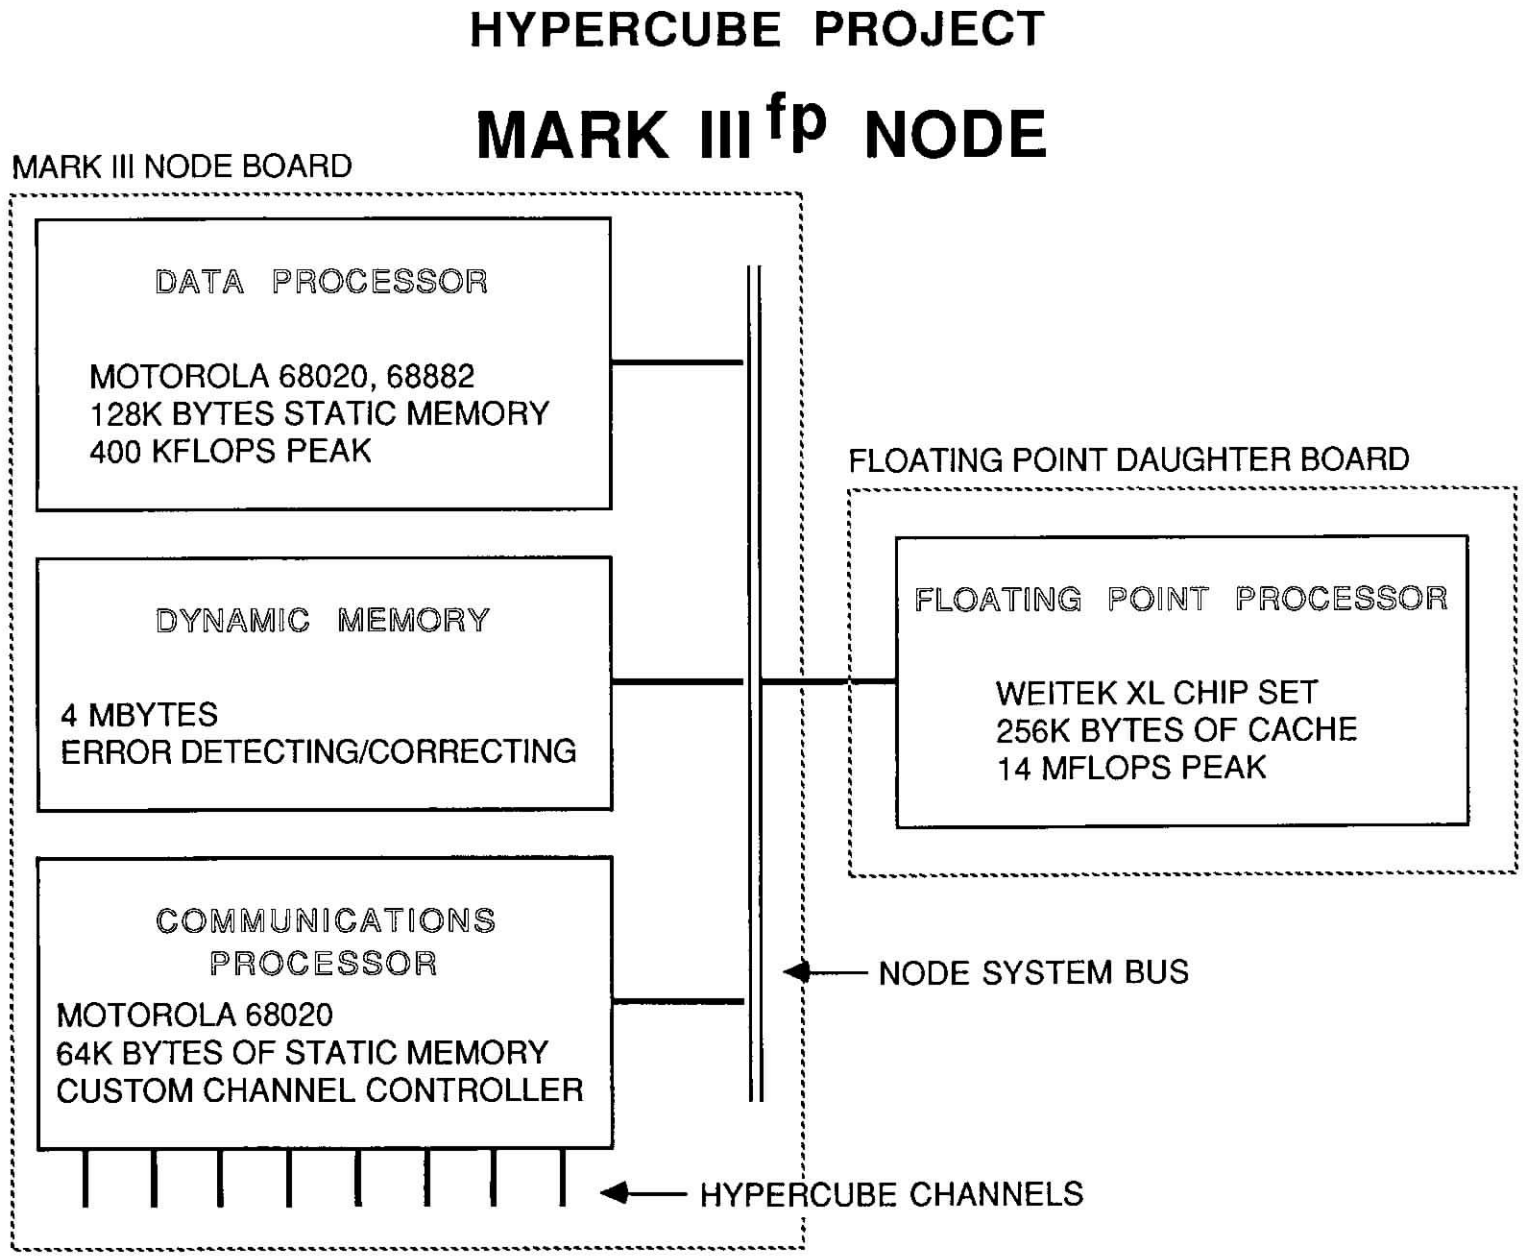

FIGURE 1a Schematic Diagram of the Concurrent I/O Subsystem & Its Connectivity to a Hypercube Architecture

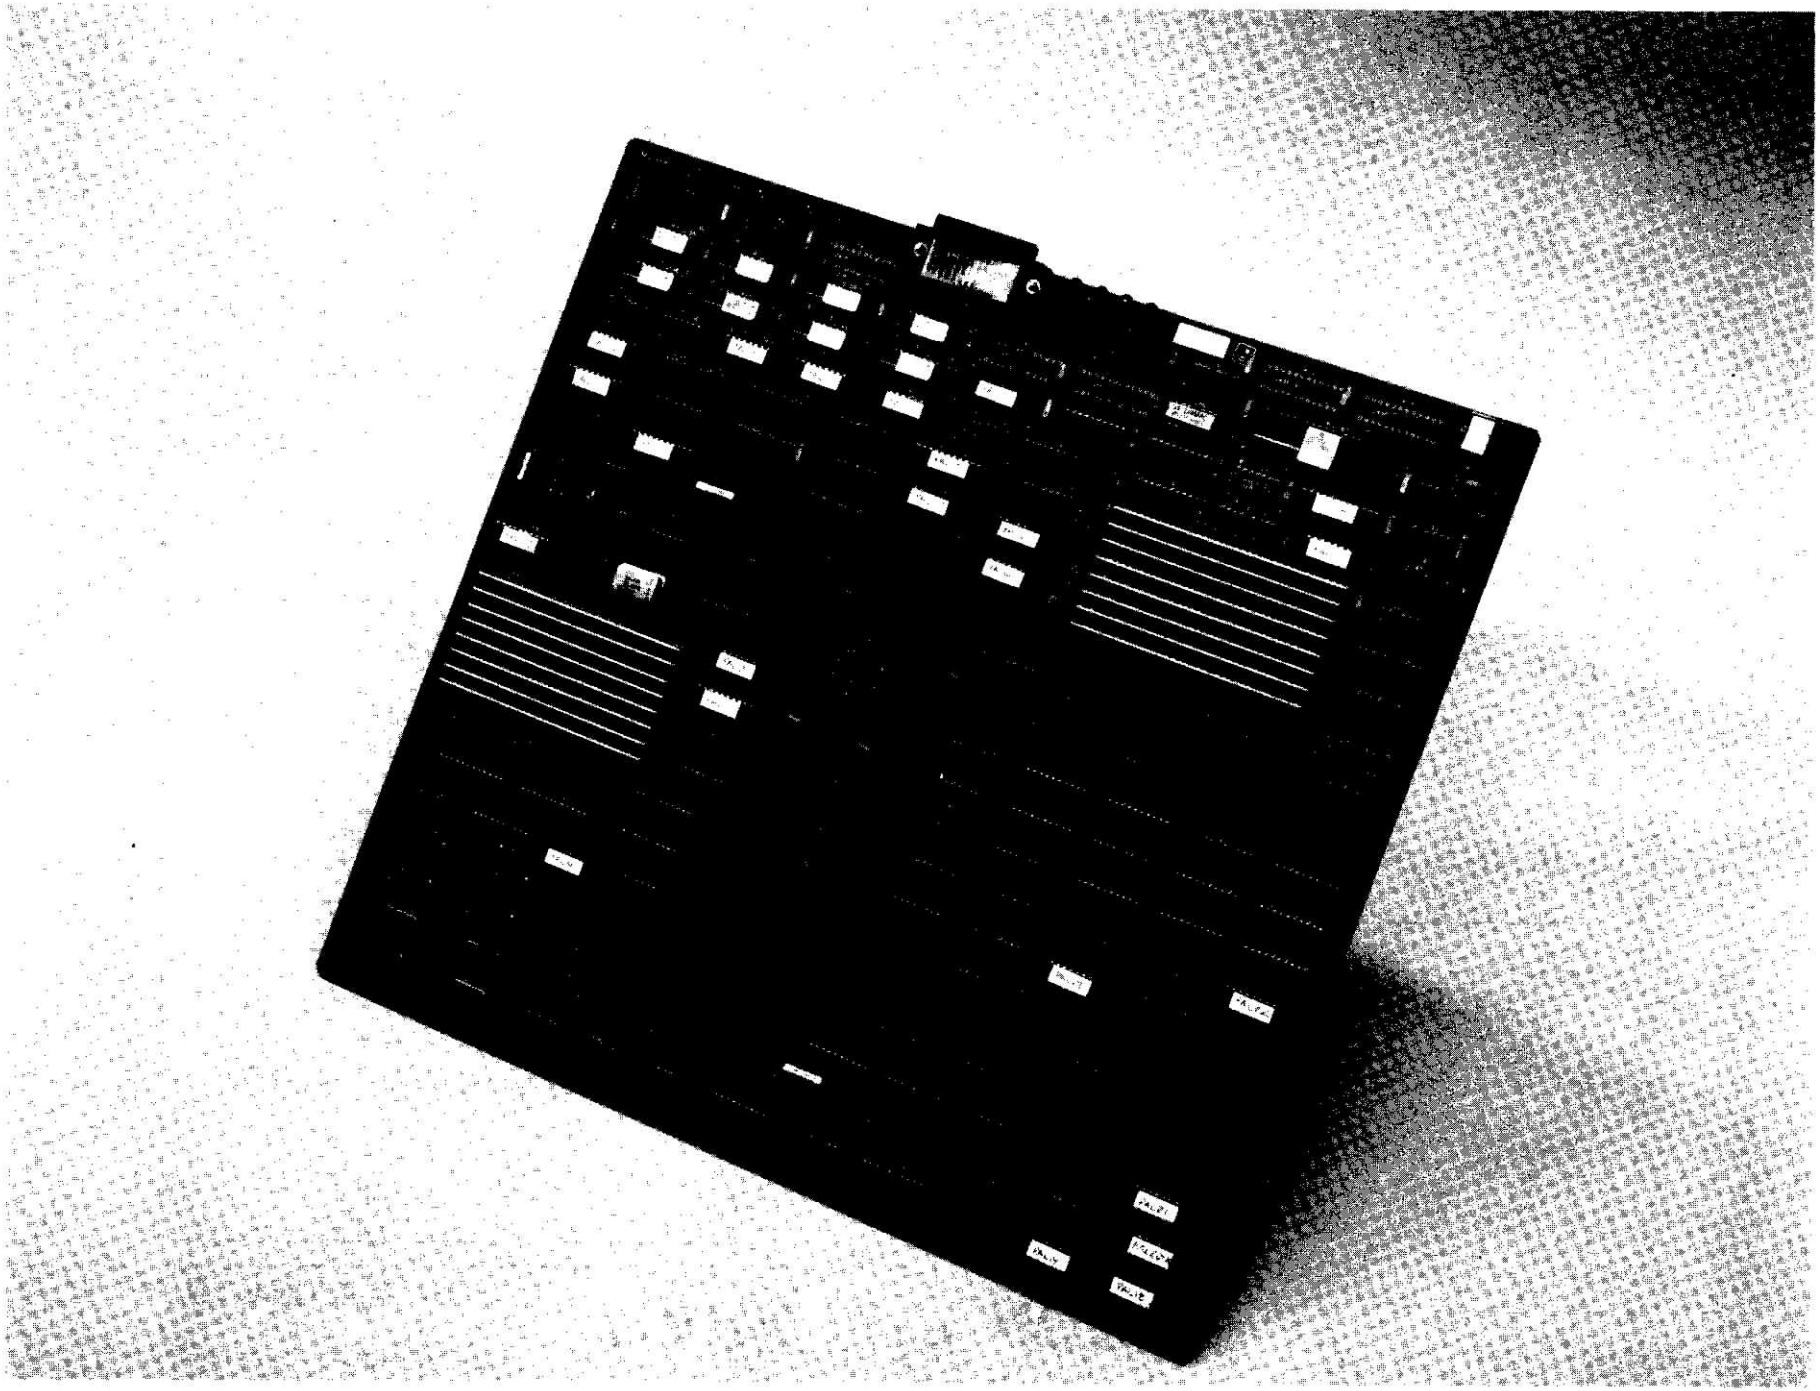

FIGURE 1b Mark III Floating Point
Daughter Board