



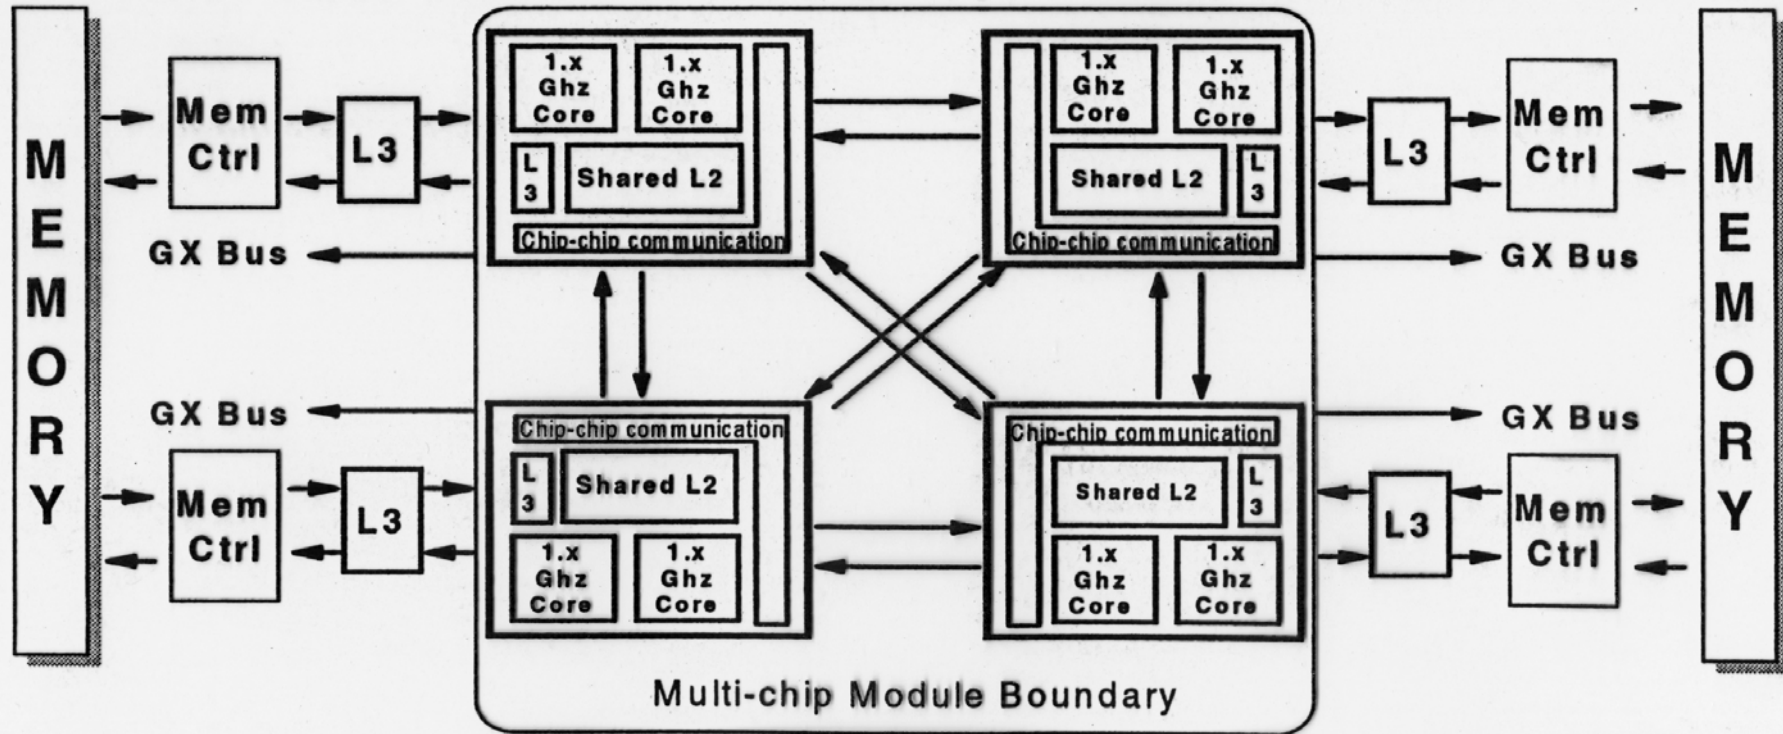
**IBM @server pSeries**

**IBM @server pSeries 690**

**Targeted for the Demands of the  
High Performance Computing  
Customer**

# pSeries 690 Multi-Chip Module (Logical View)

Four POWER4 chips (eight processors) on an MCM



## Notes:

The four GX Bus links provide connections to external I/O devices. If memory is balanced, L3 cache is shared across all processors.

Figure 2-2 The pSeries 690 Multi-Chip Module (MCM)

It is interesting to note that each POWER4 chip contains more than one mile of

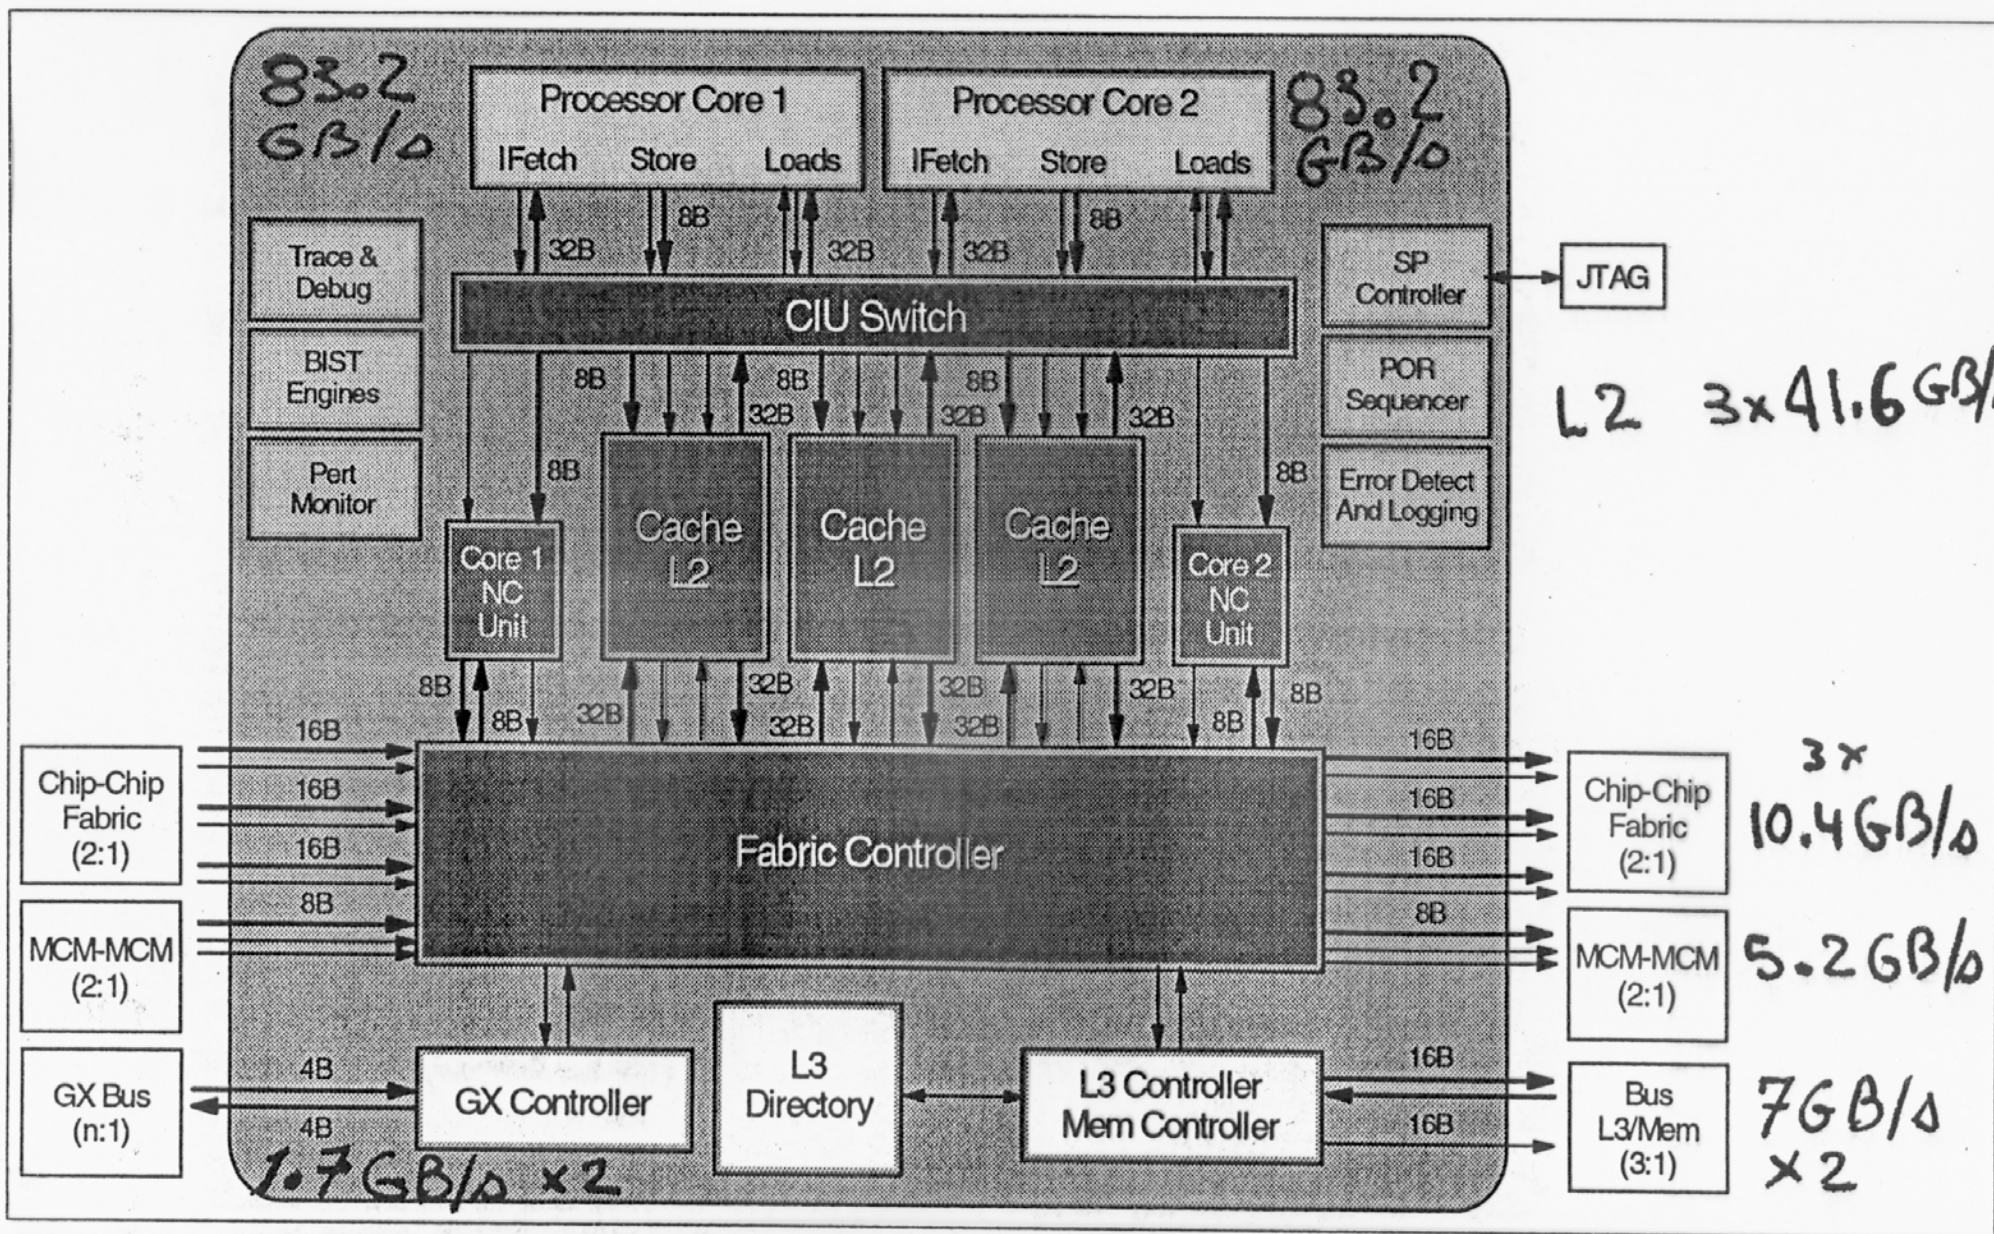


Figure 2-4 Inside the POWER4 chip

# pSeries 690 CEC (Logical View)

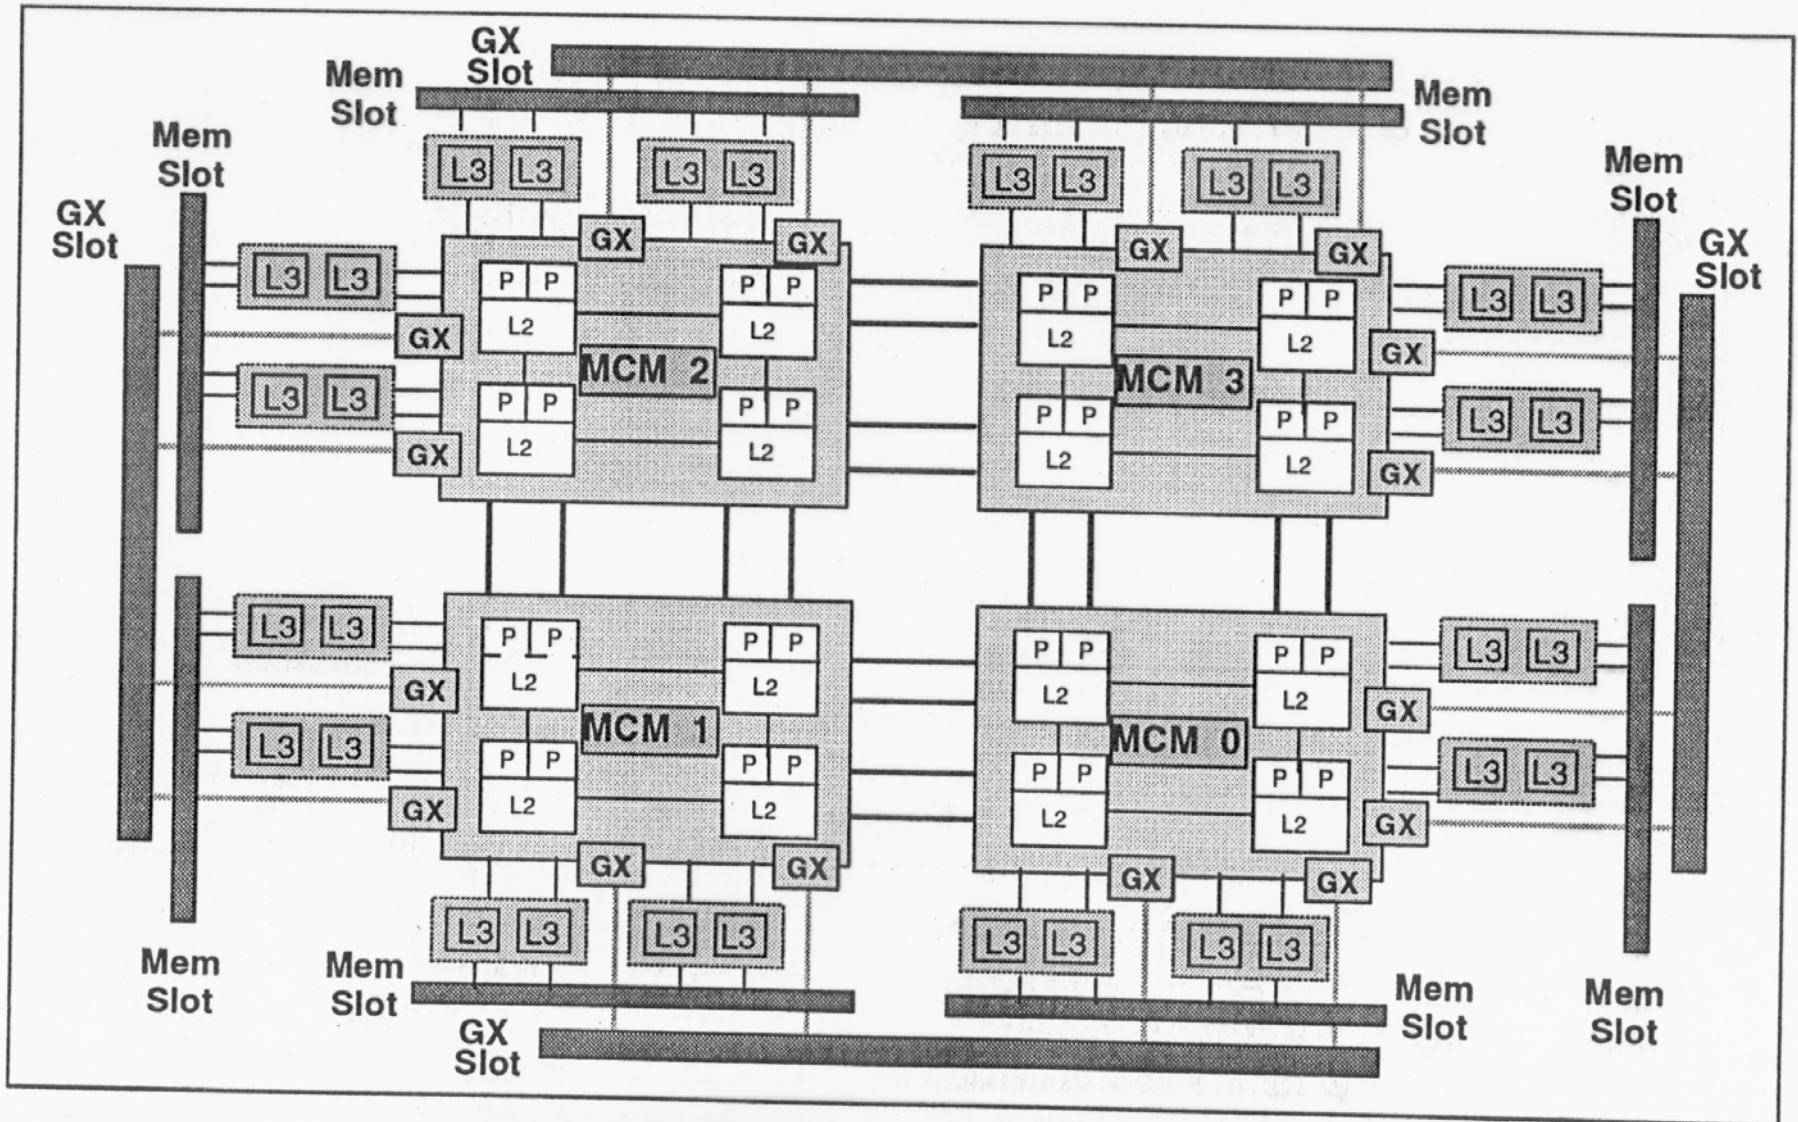
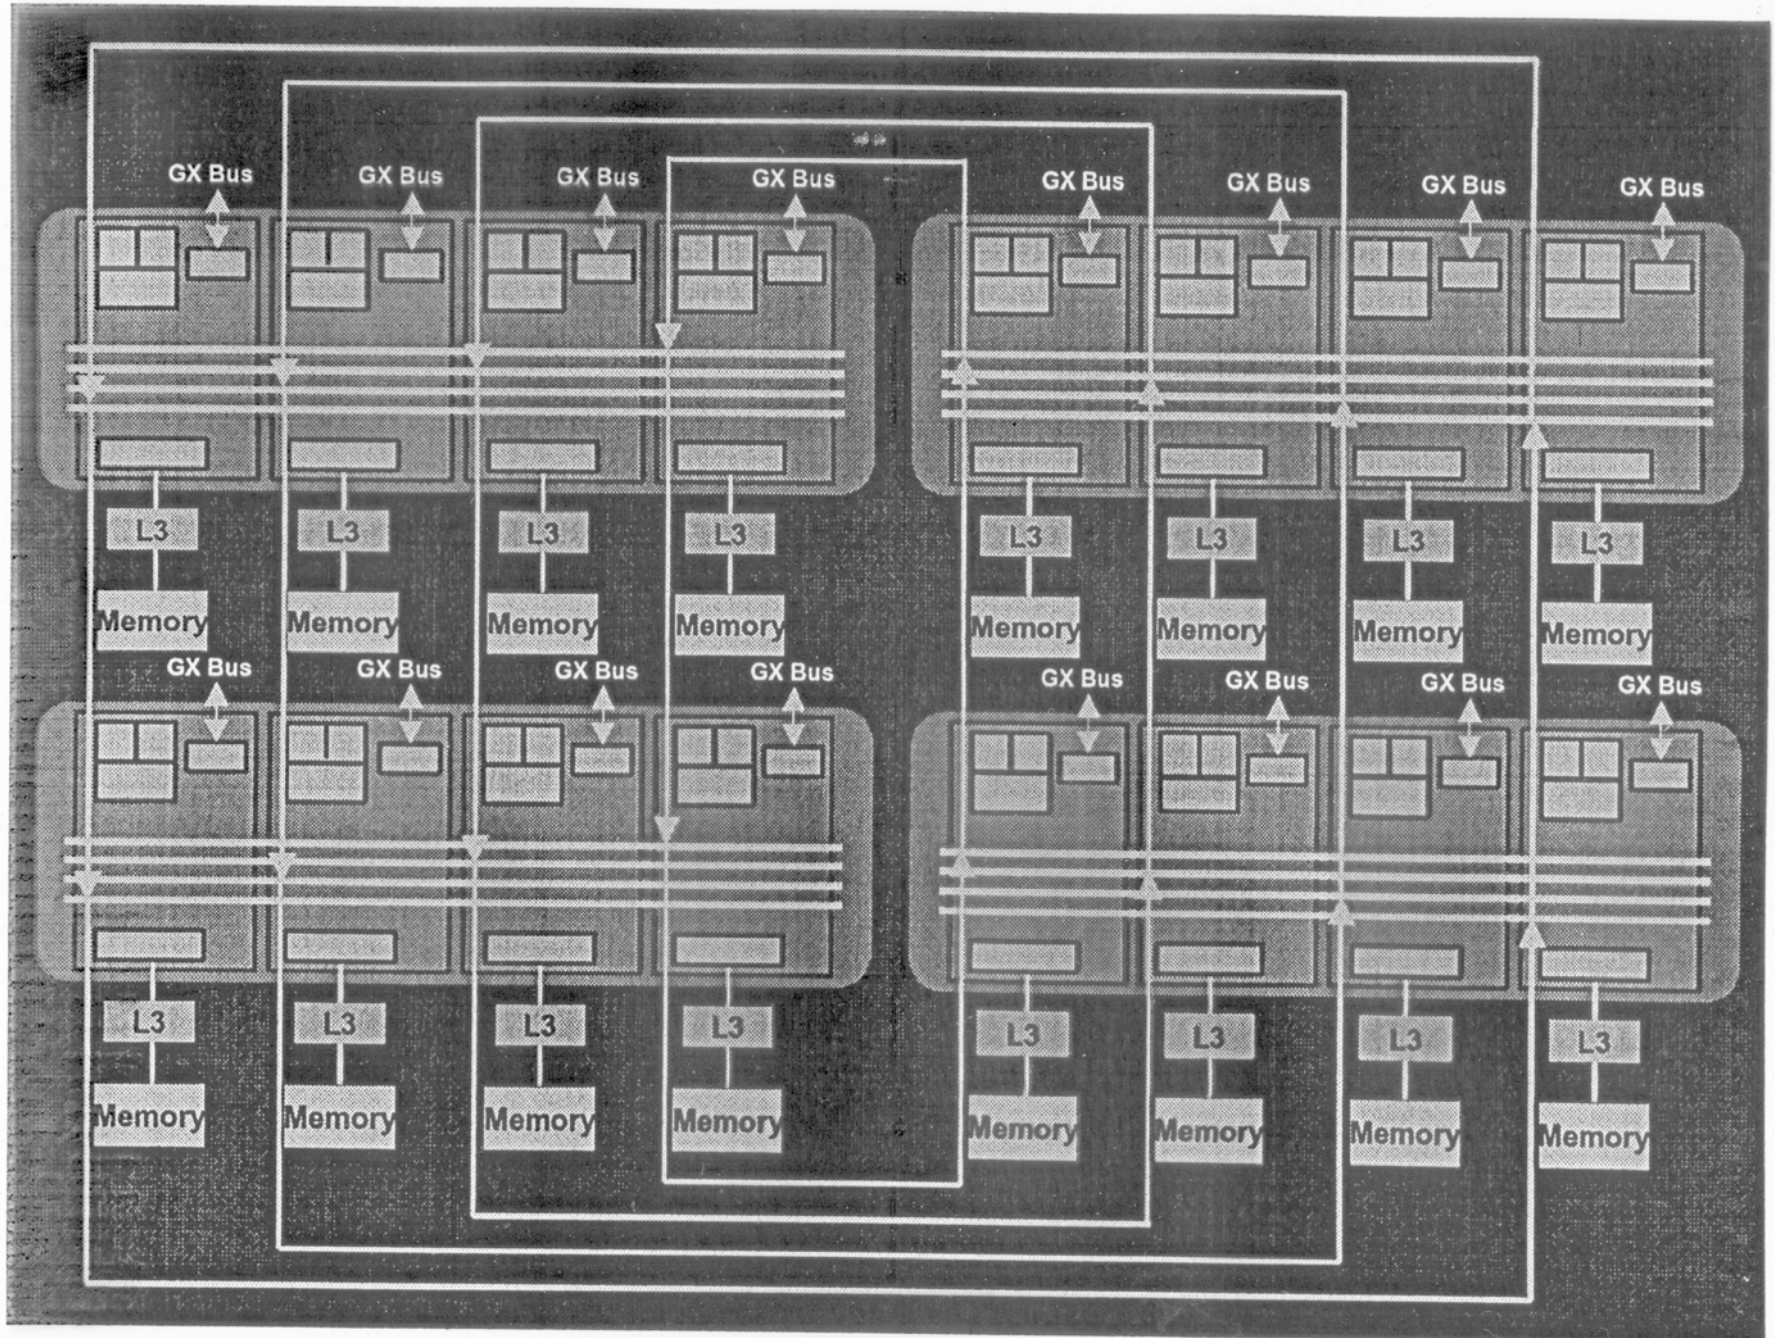



Figure 2-3 pSeries 690 fully-populated CEC

HOWEVER, IF THE L3 IS BUSY SERVING REQUESTS, IT IS NOT CACHED. INSTEAD, DATA IS NOT CACHED ON the sourcing module if it is being sourced to a chip on another module.



**Hardware Data Prefetch:** POWER4 systems employ hardware to prefetch data transparently to software into the L1 data cache. When load instructions miss sequential cache lines, either ascending or descending, the prefetch engine initiates accesses to the following cache lines before being referenced by load instructions. In order to insure the data will be in the L1 data cache, data is prefetched into the L2 from the L3 and into the L3 from memory. Figure 7 shows the sequence of prefetch operations. Eight such streams per processor are supported.

**Figure 7: POWER4 hardware data prefetch**



## Storage Hierarchy

The POWER4 storage hierarchy consists of three levels of cache and the memory subsystem. The first and second levels of the hierarchy are on board the POWER4 chip. The directory for the third level cache, the L3, is on the chip, but the actual cache is off chip. Table 3 shows capacities and organization of the various levels of the hierarchy on a per chip basis.

**Table 3: Storage hierarchy organization and size**

Component	Organization	Capacity per Chip
L1 Instruction Cache	Direct map, 128-byte line managed as 4 32-byte sectors	128 KB (64 KB per processor)
L1 Data Cache	2-way, 128-byte line	64 KB (32 KB per processor)
L2	8-way, 128-byte line	~ 1.5 MB
L3	8-way, 512-byte line managed as 4 128-byte sectors	32 MB
Memory	---	0-16 GB

**Table 2: Rename resources**

Resource Type	Logical Size	Physical Size
GPRs	32 (36)	80
FPRs	32	72
CRs	8 (9) 4-bit fields	32
Link/Count Registers	2	16
FPSCR	1	20
XER	4 fields	24

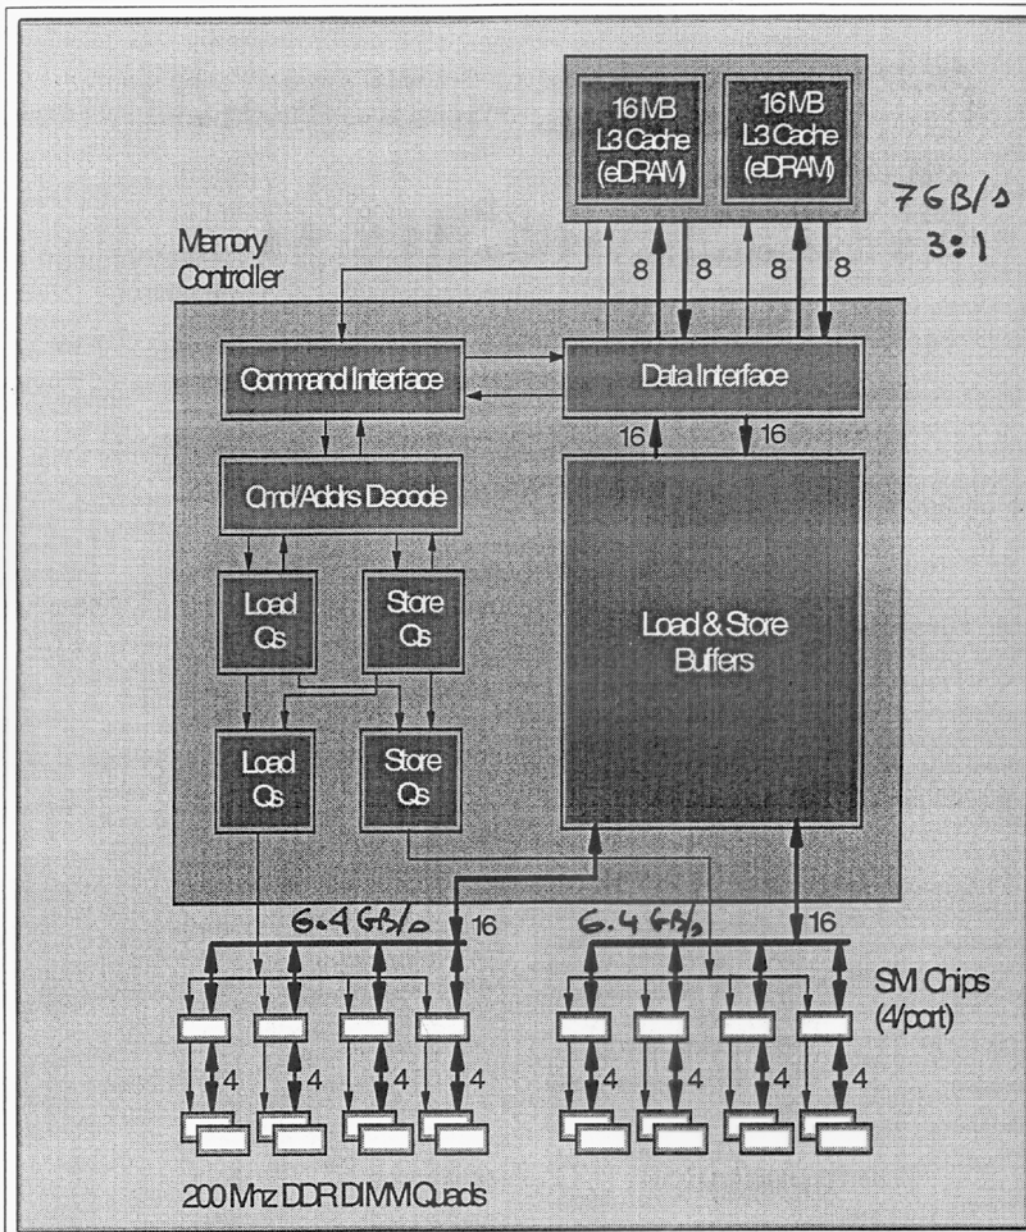


Figure 2-5 Logical view of POWER4 memory

### Memory balancing

In general, a balanced memory configuration is critical for achieving optimum and consistent performance in both commercial and technical computing environments. To maximize memory performance on the pSeries 690, memory interleaving is employed. If an MCM has two memory cards of the same size



## 2.2.2 The memory subsystem

The dividing line for where the memory subsystem begins is not black and white. For the purposes of this book, we define the memory storage hierarchy as follows:

1. L1 data and instruction caches
2. L2 shared caches
3. L3 directories and controllers
4. L3 off-chip caches
5. Memory controllers and cards (packaged within memory books)

Table 2-4 lists the organization and per-chip capacity for each of the components in the memory storage hierarchy.

Table 2-4 POWER4 memory and cache organization and per-chip capacity

Component	Organization	Capacity (per chip)
L1 Instruction	Direct map 128-byte line (4 x 32-byte sectors)	128 KB/chip (64 KB/processor)
L1 Data	Two-way 128-byte line	64 KB/chip (32 KB/processor)
L2	Eight-way 128-byte line	1440 KB/chip
L3	Eight-way 512-byte line (4 x 128-byte sectors)	32 MB/chip (128 MB/MCM)
Memory	N/A	0 - 16 GB

The following brief descriptions of the caches are excerpted from *The POWER4 Processor Introduction and Tuning Guide*, SG24-7041.

### L1 instruction cache

Each POWER4 microprocessor (which is one of the components on the POWER4 chip) contains an L1 instruction cache that is 64 KB in size, direct mapped, and indexed by the effective address of the instruction cache line. It is capable of either one 32-byte read or write each cycle.

Figure 10: I/O logical view

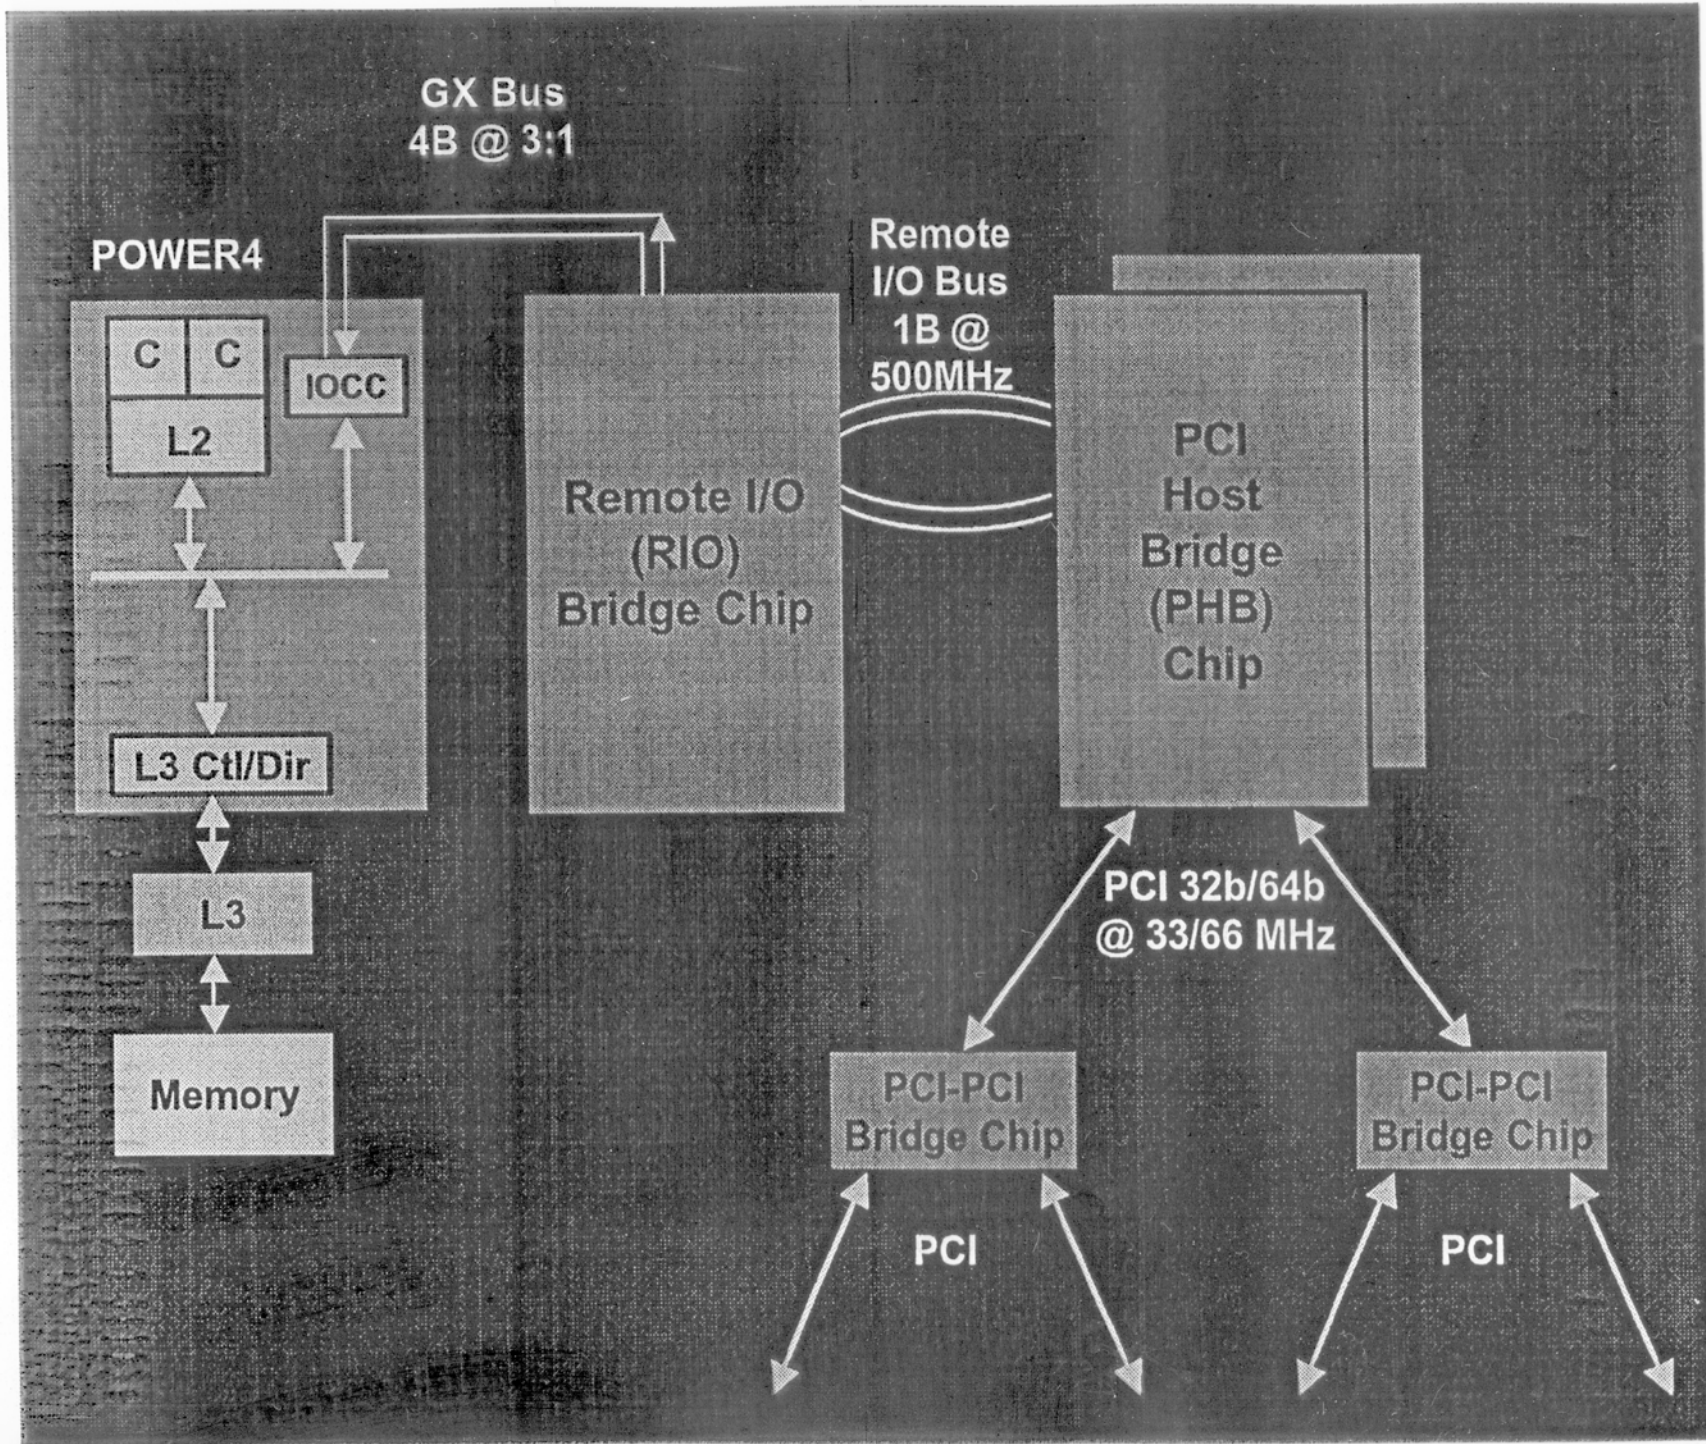


Figure 8: POWER4 multi-chip module with four chips

