Einführung in das System z Mainframe

Prof. Dr.-Ing. Wilhelm G. Spruth

2. Juli 2007
Humboldt Universität Berlin

Teil 4

System z Hardware

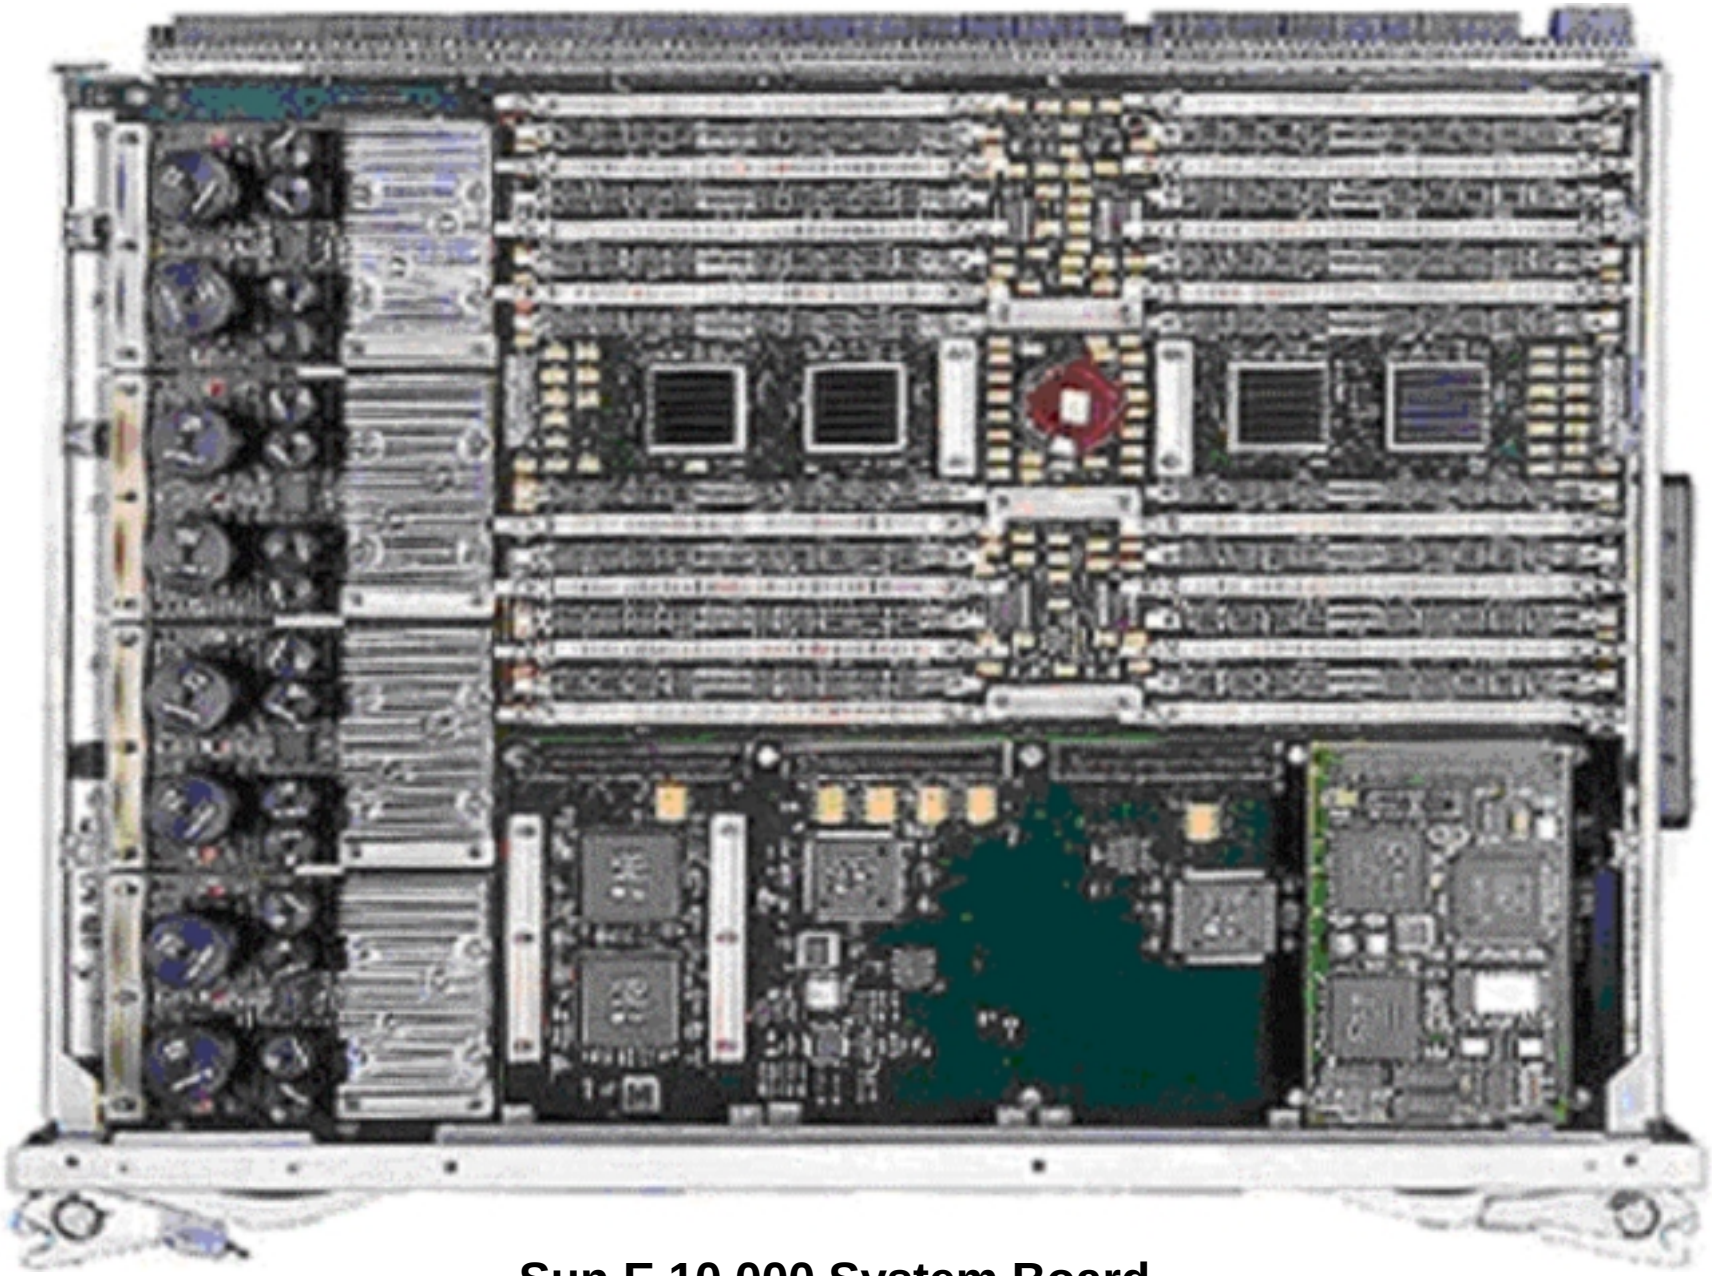

Sun E 10 000 System Board

HP Superdome Cell Board

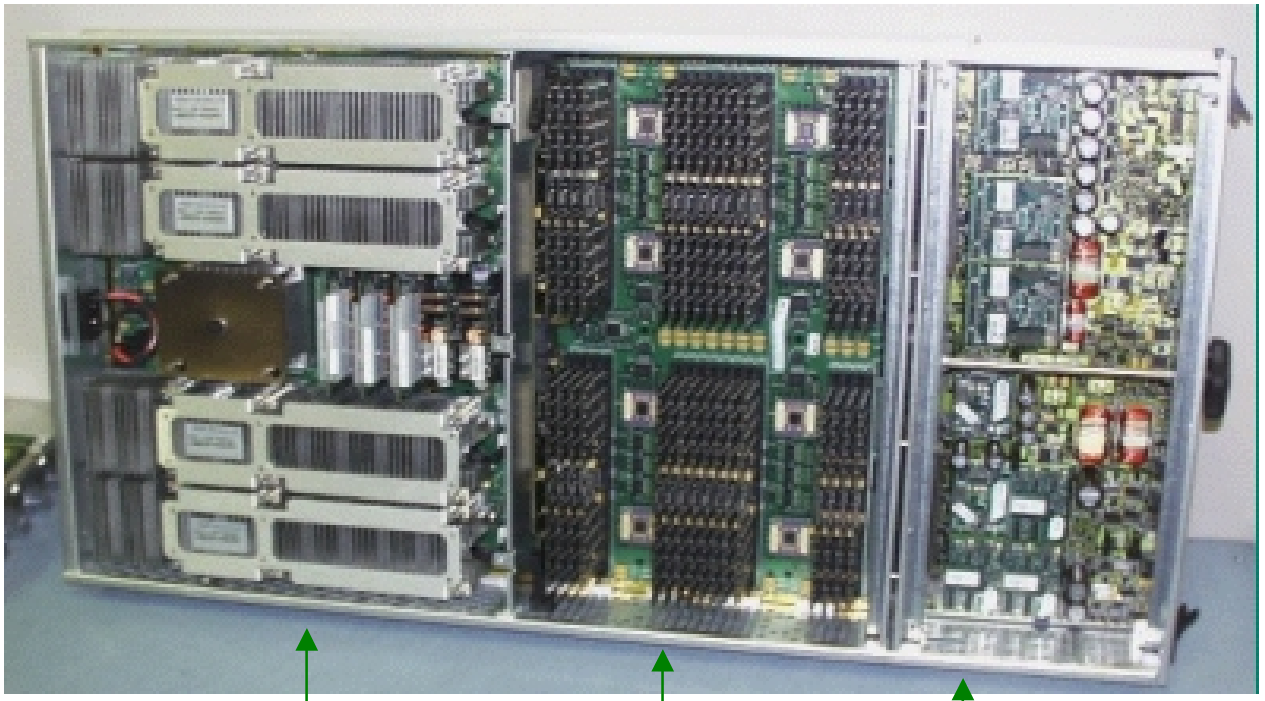

4 Itanium 2 CPUs
1,5 GHz

32 Gbyte
Hauptspeicher

E/A Bus
Anschlüsse
12 PCI

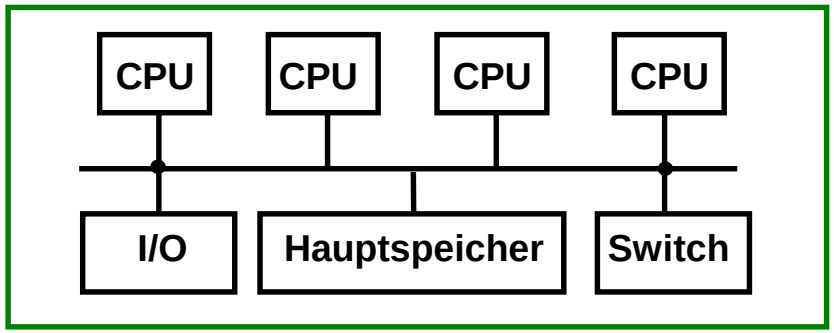

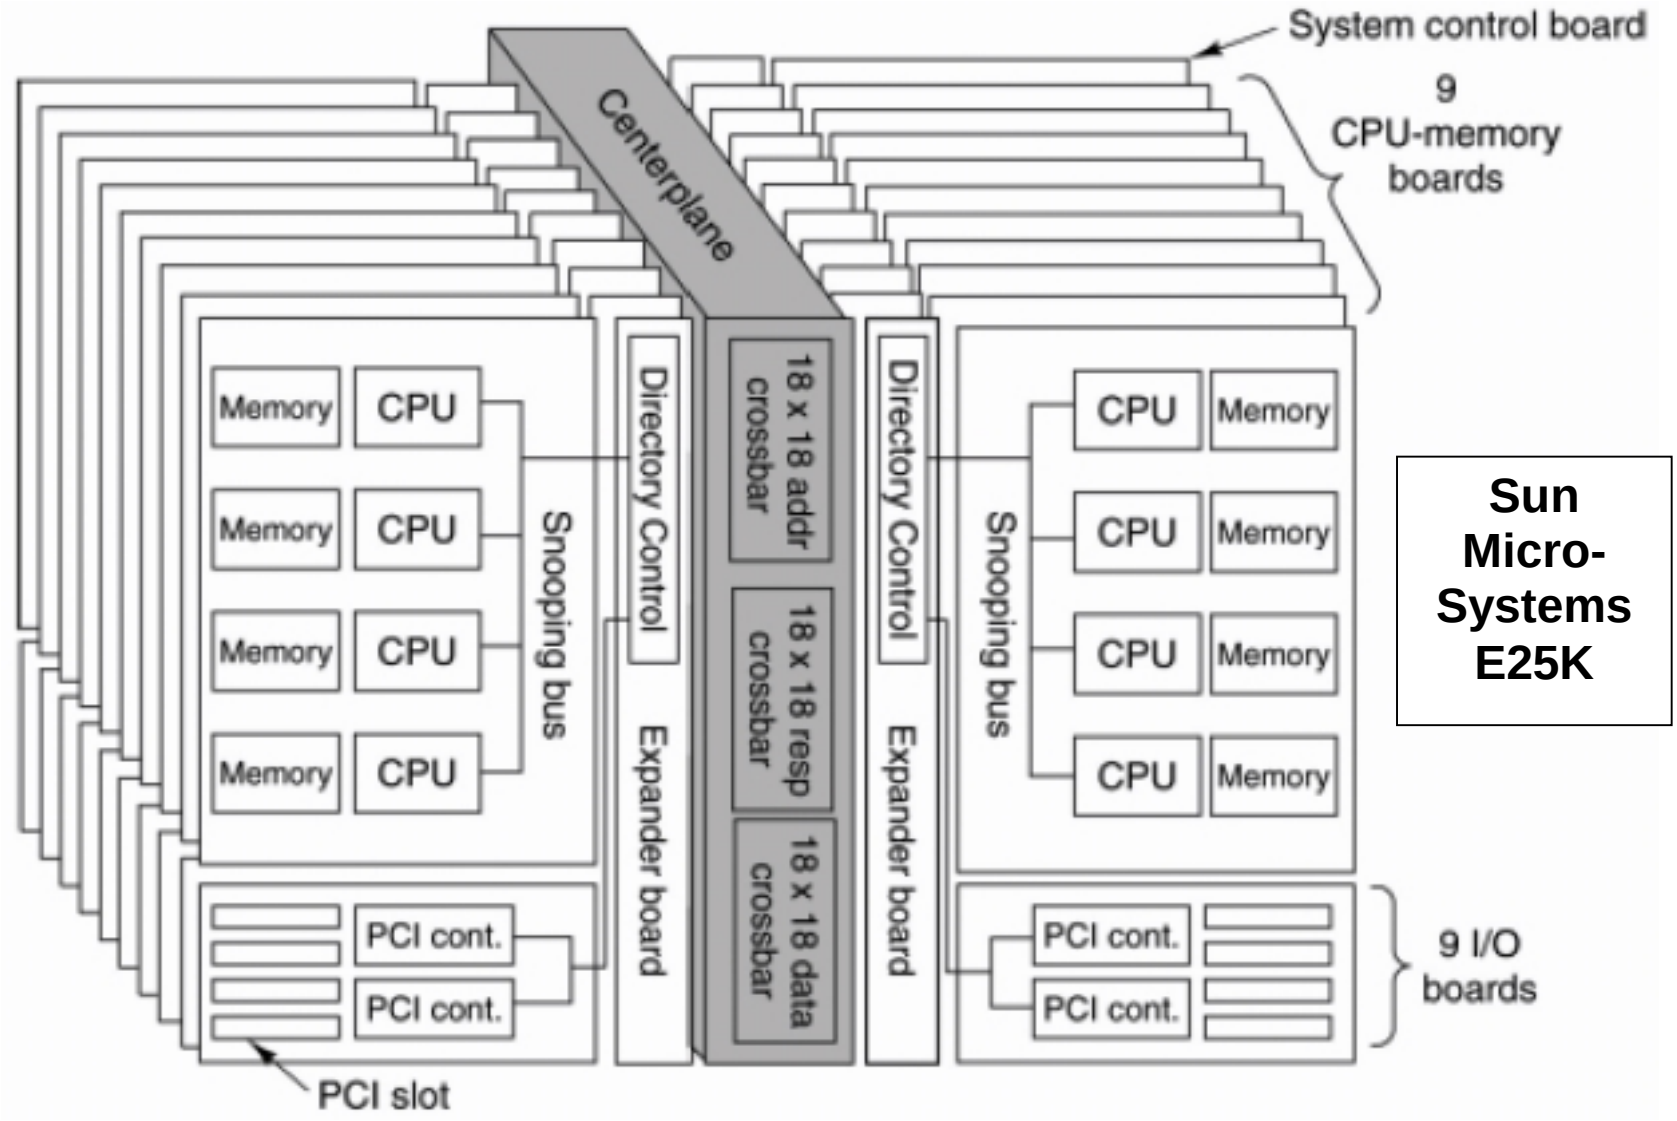

Sun Micro-Systems E25K

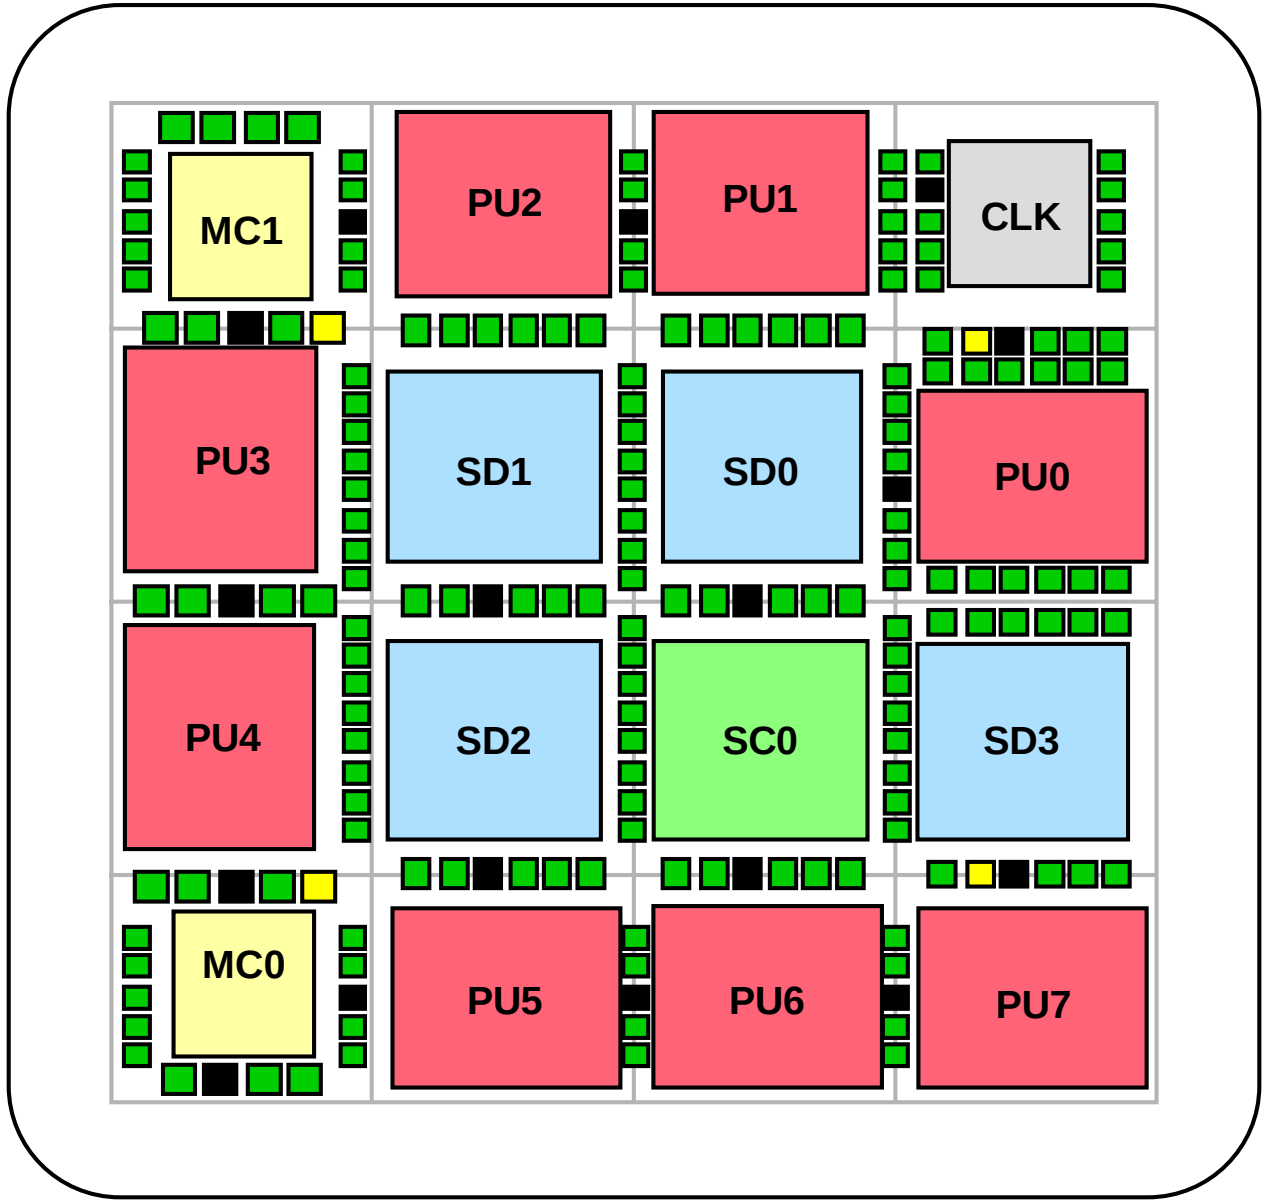
MC1
PU2
PU1
CLK
PU3
SD1
SD0
PU0
PU4
SD2
SC0
SD3
MC0
PU5
PU6
PU7

z9 Modell S54

- **Advanced 95mm x 95mm MCM**
  - –102 Glass Ceramic layers
  - –16 chip sites, 217 capacitors
  - –0.476 km of internal wire

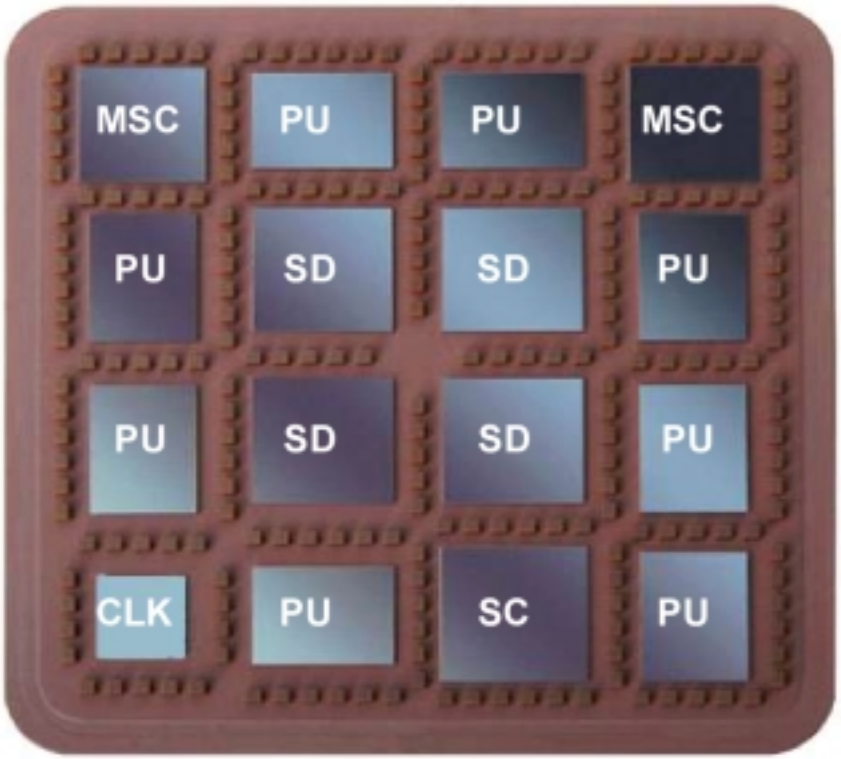

- **CMOS 10K chip Technology**
  - – Copper interconnections, 10 copper layers
  - – 8 PU chips/MCM
    - •121 million transistors/chip
    - •L1 cache/PU (256 KB I-cache,256 KB D-cache)
    - •0.58 ns Cycle Time
  - – 4 System Data (SD) cache chips/MCM
    - •40 MB L2 cache per Book
      - – 660 million transistors/chip
  - – One Storage Control (SC) chip
    - •162 million transistors
    - •L2 cache crosspoint switch
    - •L2 access rings to/from other MCMs
  - – Two Memory Storage Control (MSC) chips
    - •24 million transistors/chip
    - •Memory cards (L3) interface to L2
    - •L2 access to/from MBAs (off MCM)
  - – One Clock (CLK) chip - CMOS 8S
    - •Clock and ETR Receiver

Zwei CPU Cores auf einem Chip

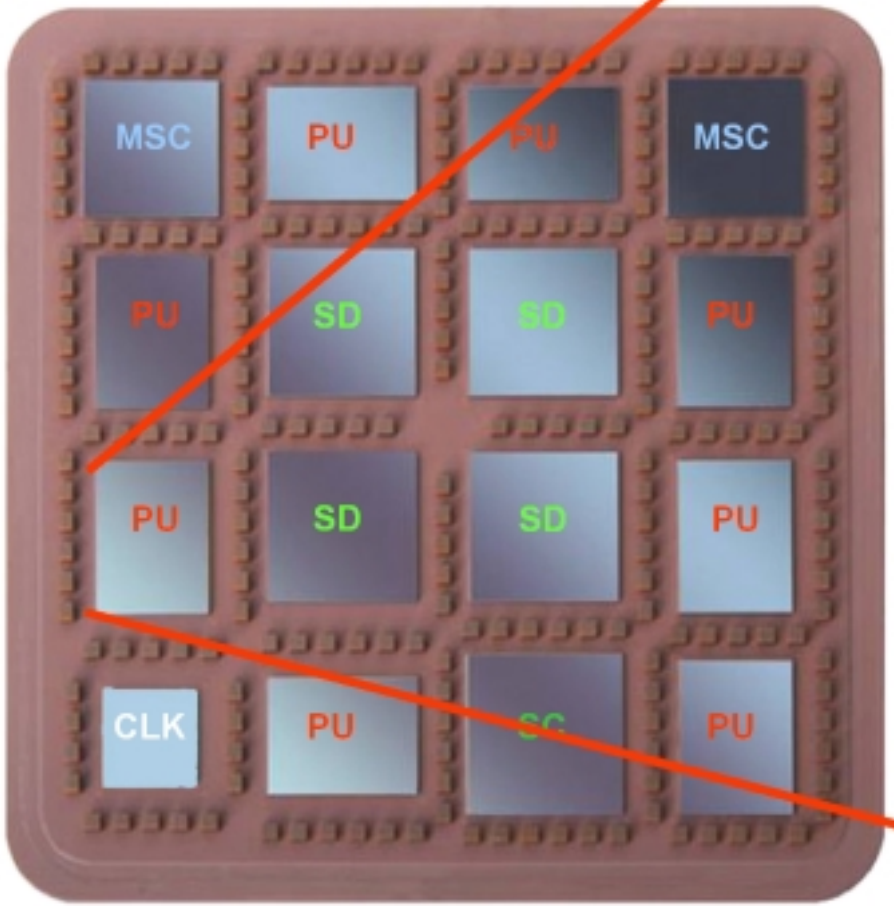

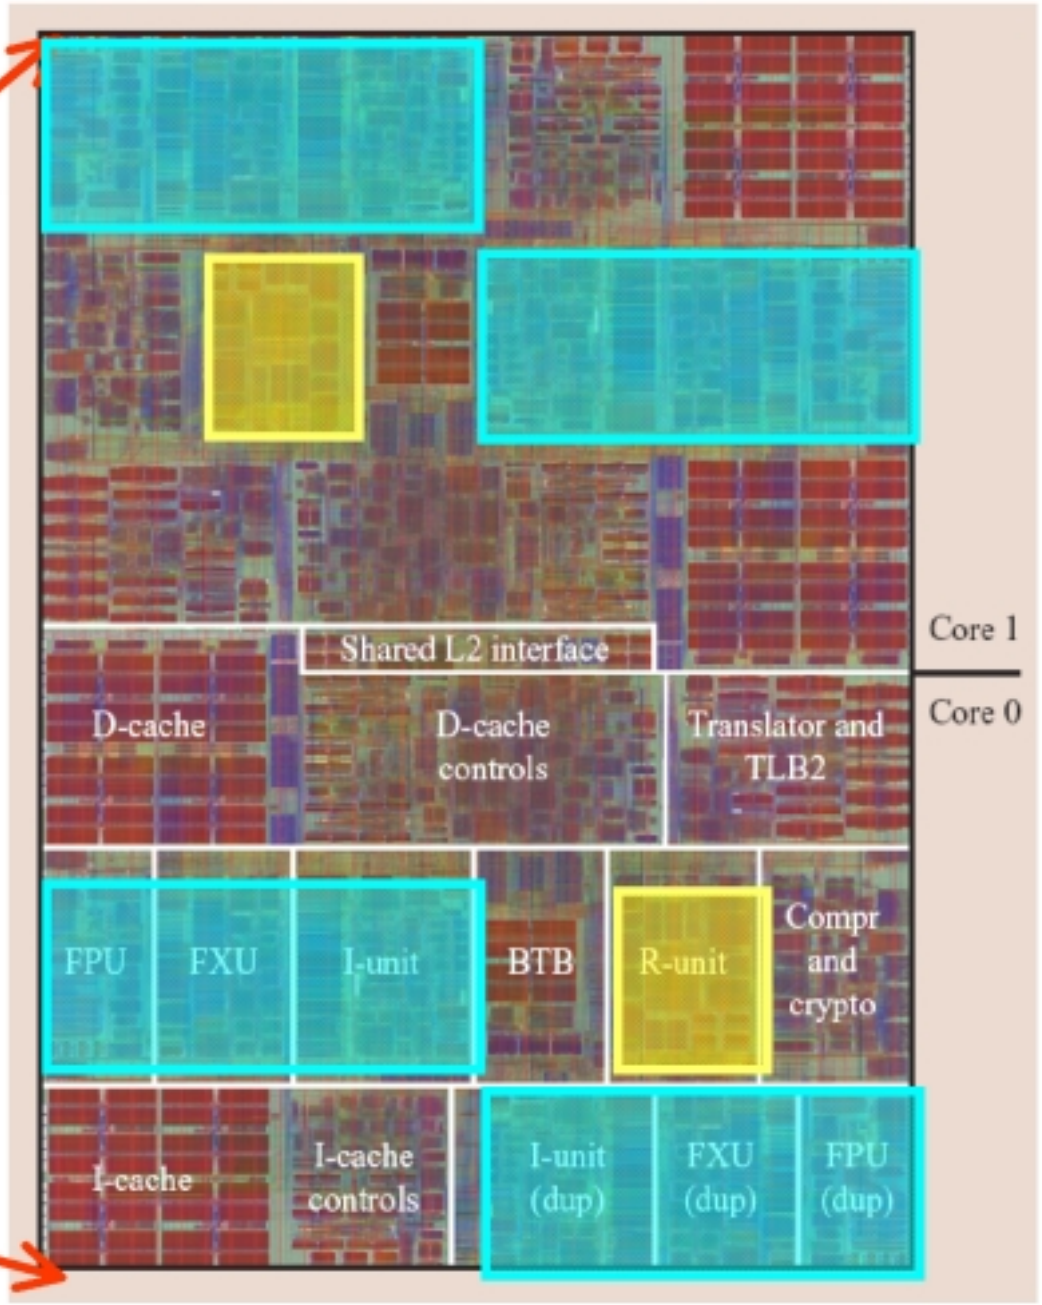

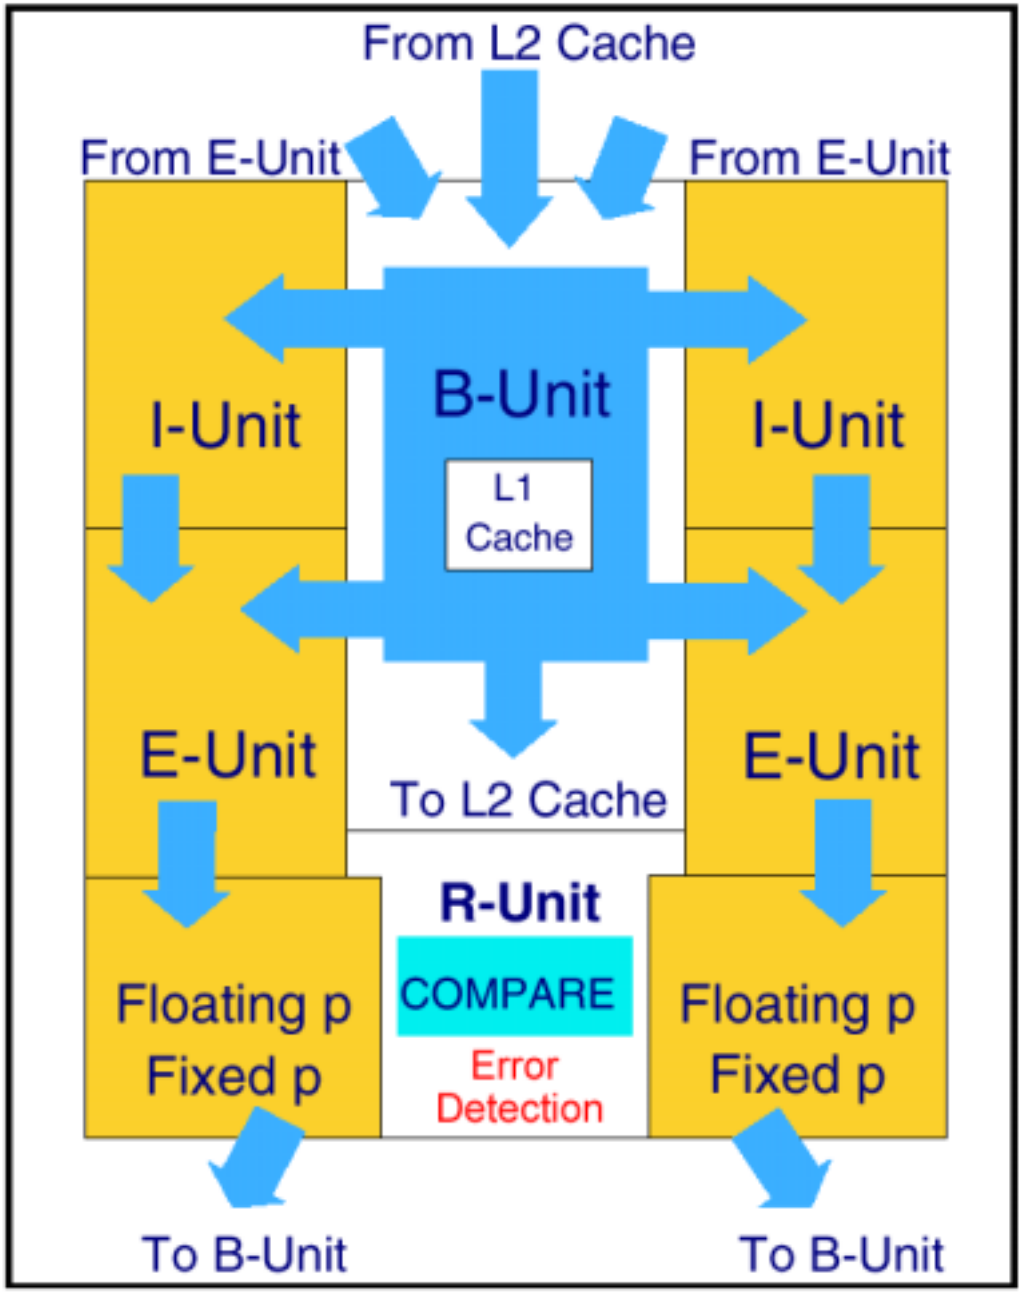

- Processing Unit (PU)
  - Dual processor
  - I-Unit
  - E-Unit
  - Floating Point function
  - Simple yet complete error detection mechanism
  - Data flow - parity checked
  - Address paths - parity checked
  - L1 Cache - parity checked
  - Processor logic (I - E - F) -
    Duplicated, then compared output.
    Error detection for mis-compare

z9 Modell S54

90 nm Prozess

Uniprozessor 600 MIPS, Gesamtleistung 18 000 MIPS

512 GByte Hauptspeicher

60 LPAR

1,2 x 109 $ Entwicklungskosten

> 6 Mill. LOC Firmware

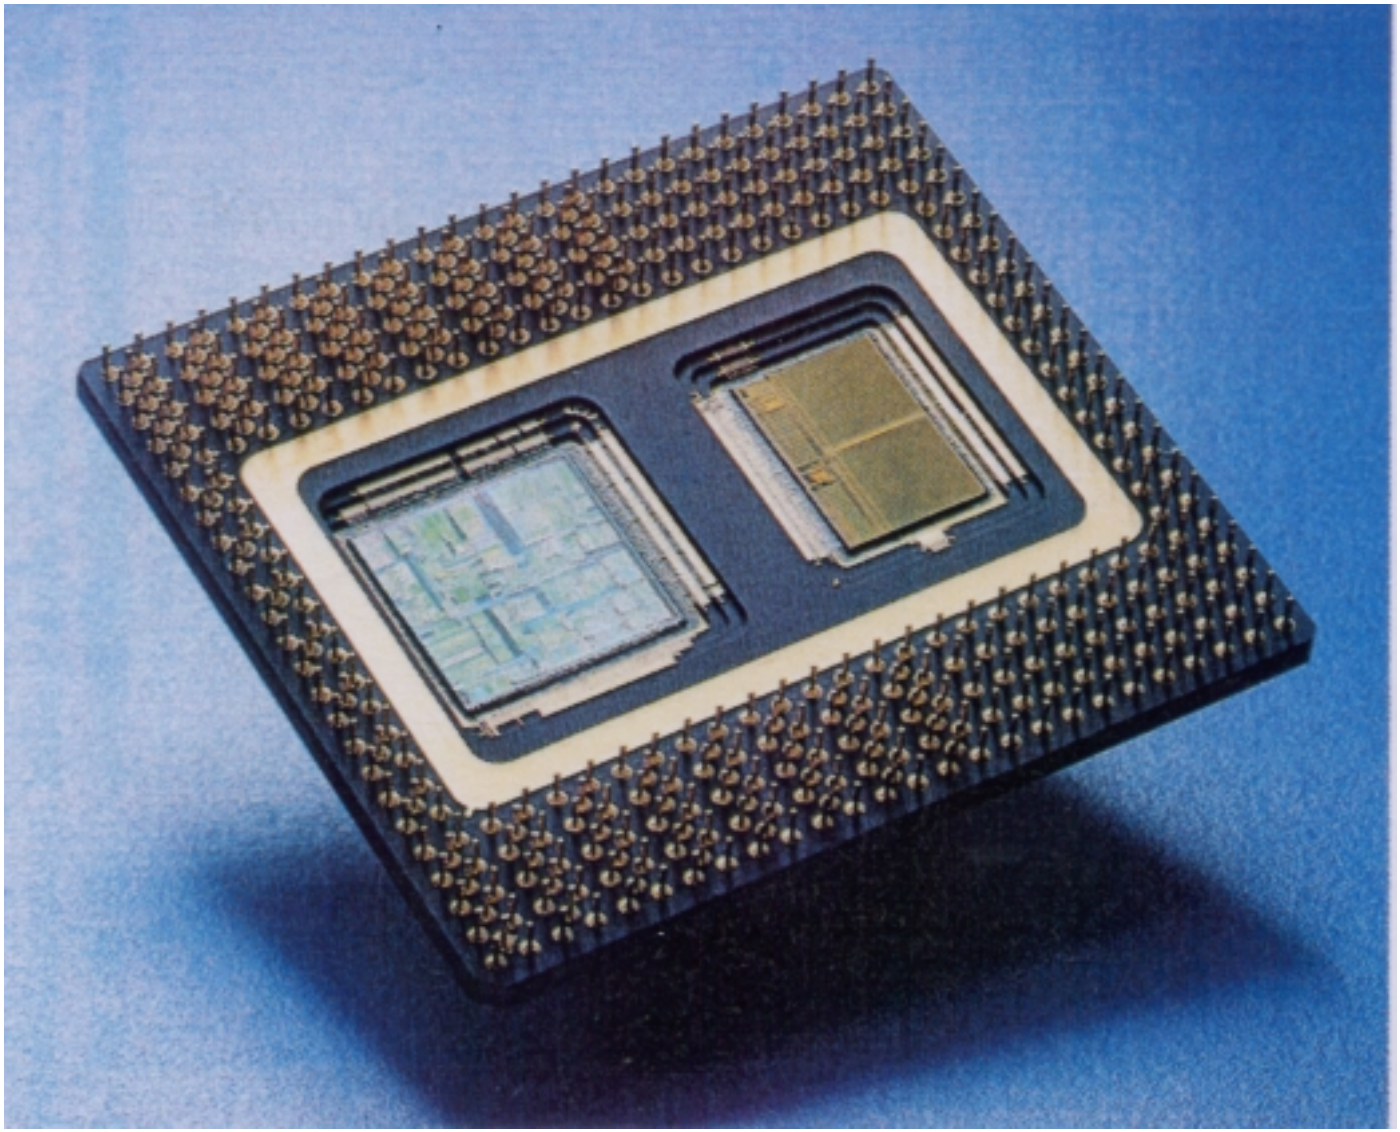

Pentium Pro

387 Pin Multi Layer Ceramic (MLC) Multi Chip Carrier (MCM) Module

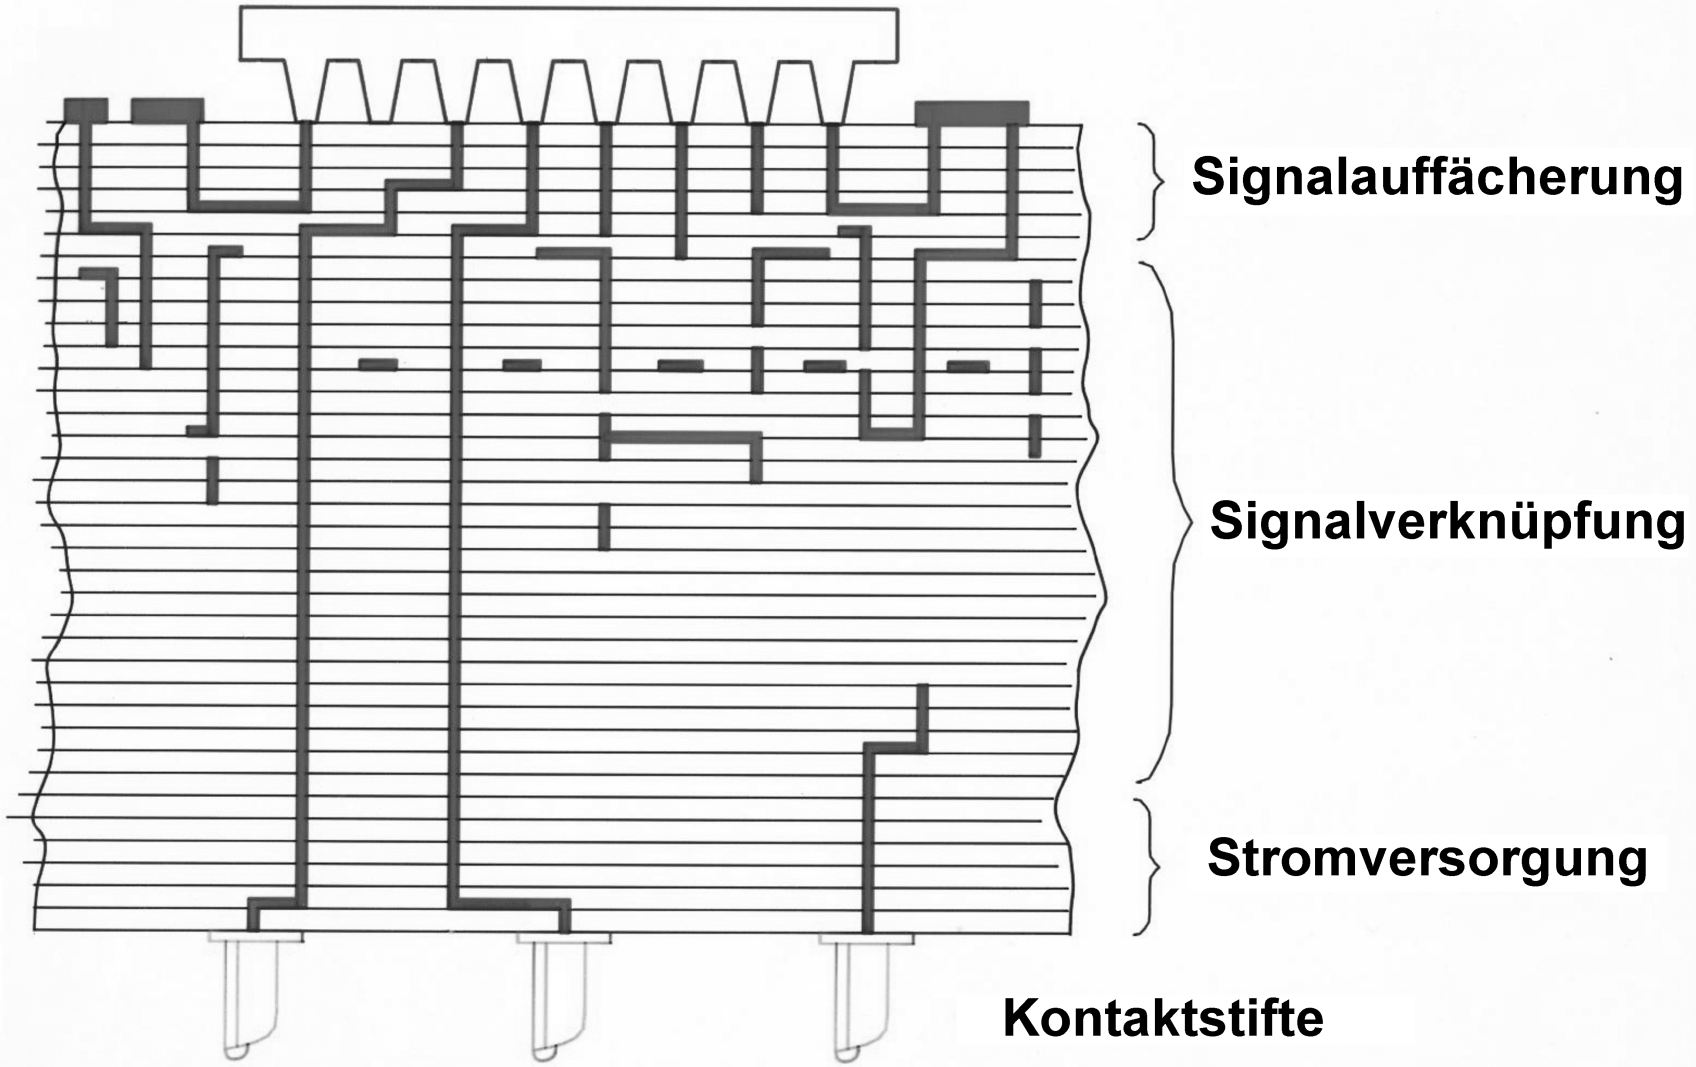

101 + 6 Schichten, über 4000 Kontaktstifte
1 km Verdrahtungslämge

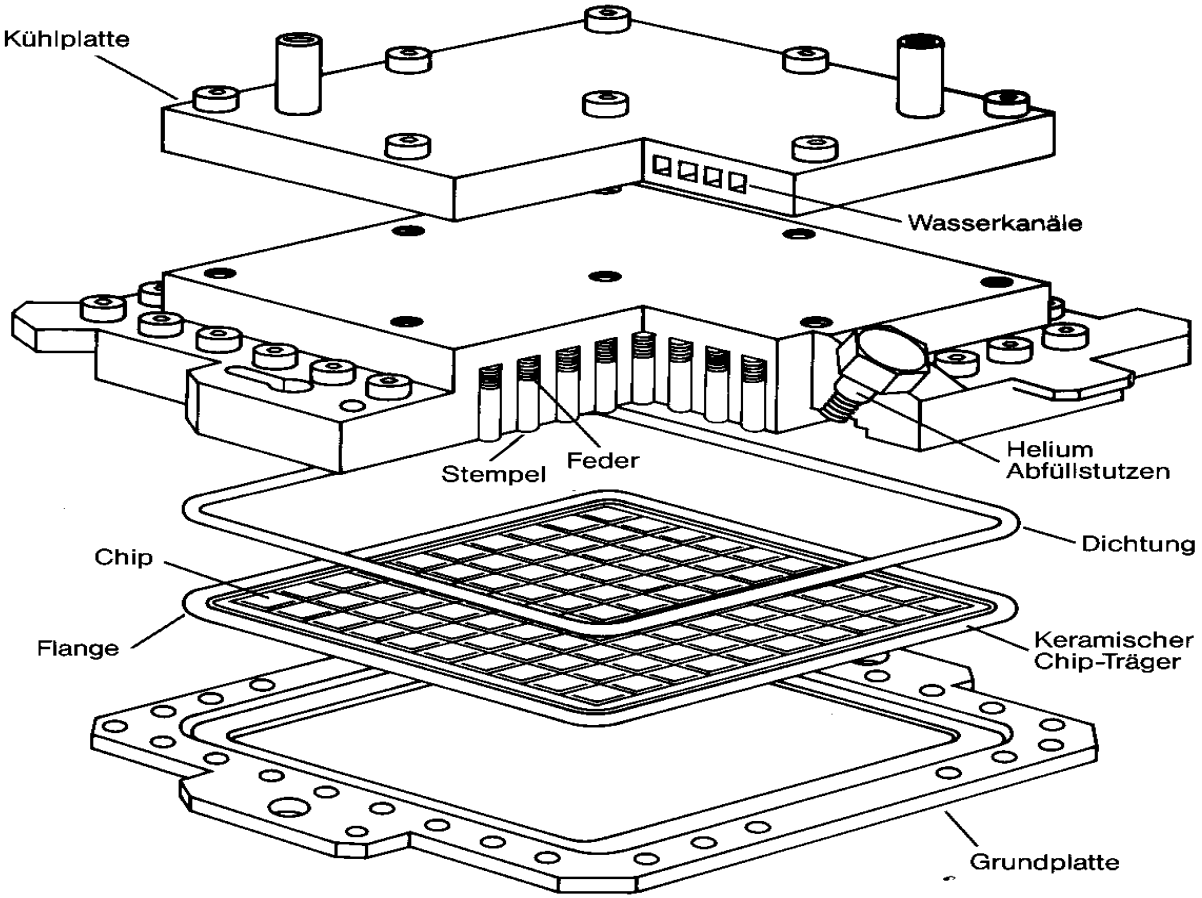

Aufbau eines „Thermal Conduction Module“

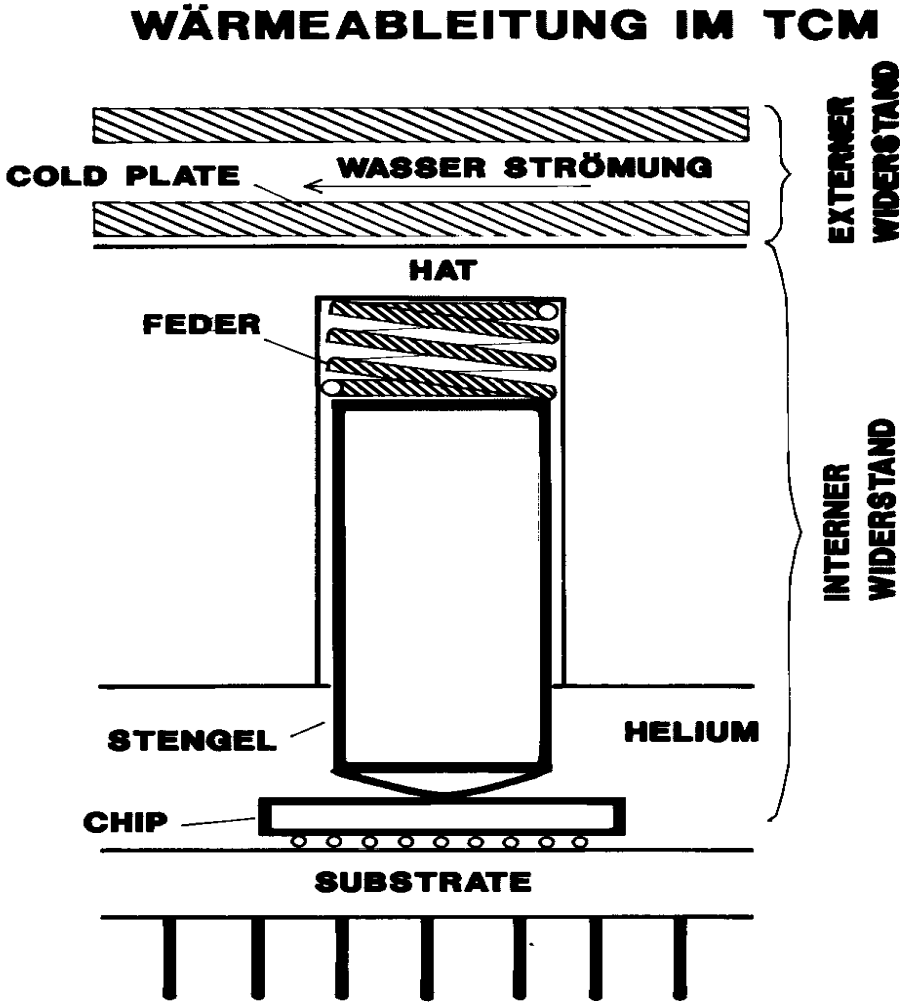

TCM Wärmeübergang

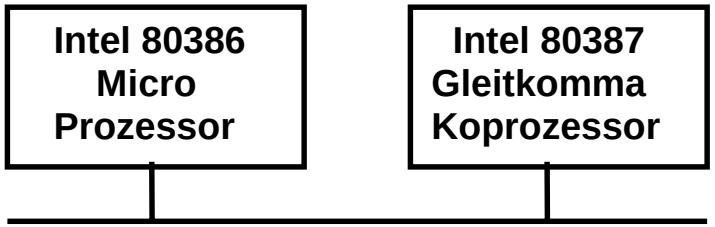

Ohne 80387 verursacht die Ausführung eines Gleitkomma-Maschinenbefehls eine Programmunterbrechung. Die Unterbrechungsroutine führt die Gleitkommaoperation aus.

Intel 80486 Micro-Prozessor

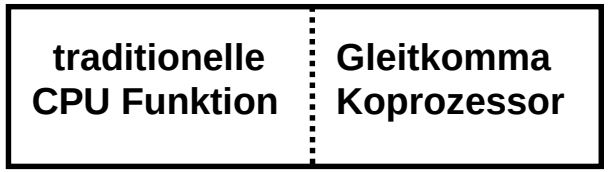

Beim Intel 808486 und Pentium/AMD ist die Gleitkommafunktion in das Micro-Prozessor Chip integriert

IBM z990 Micro-Prozessor

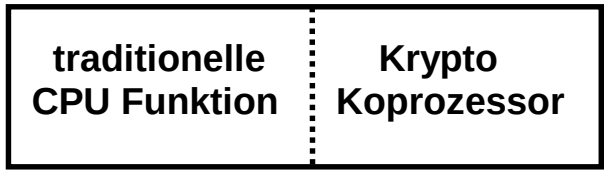

Bei den IBM z9, z990 und 890 Micro-Prozessoren ist die Krypto-Koprozessor Funktion in das Micro-Prozessor Chip integriert.
Ca. 150 zusätzliche Maschinenbefehle.

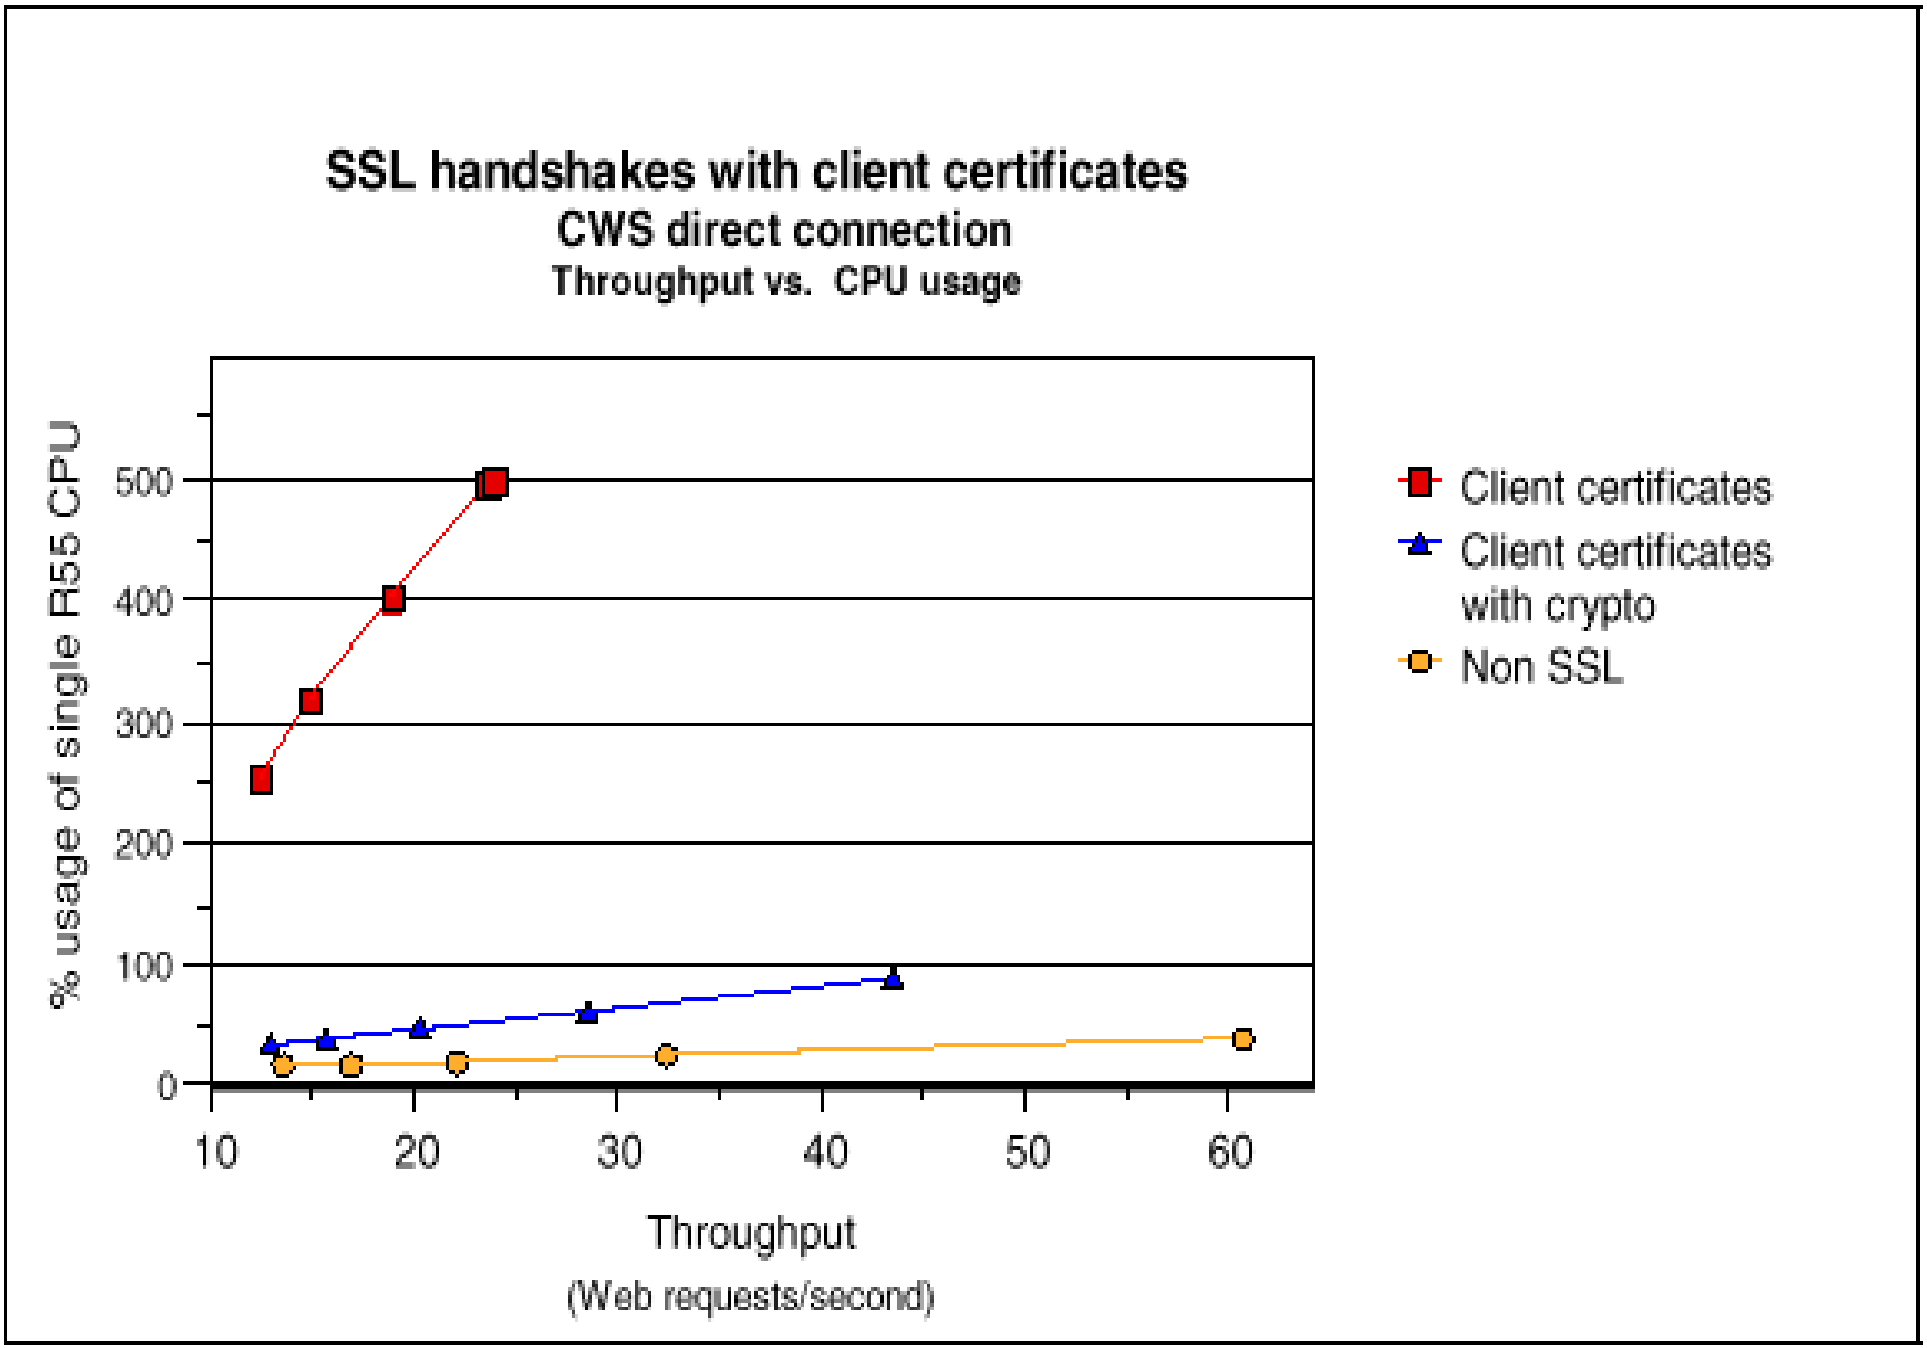
SSL handshakes with client certificates
CWS direct connection
Throughput vs. CPU usage
% usage of single R55 CPU
500
400
300
200
100
0
10
20
30
40
50
60
Throughput
(Web requests/second)
Client certificates
Client certificates with crypto
Non SSL

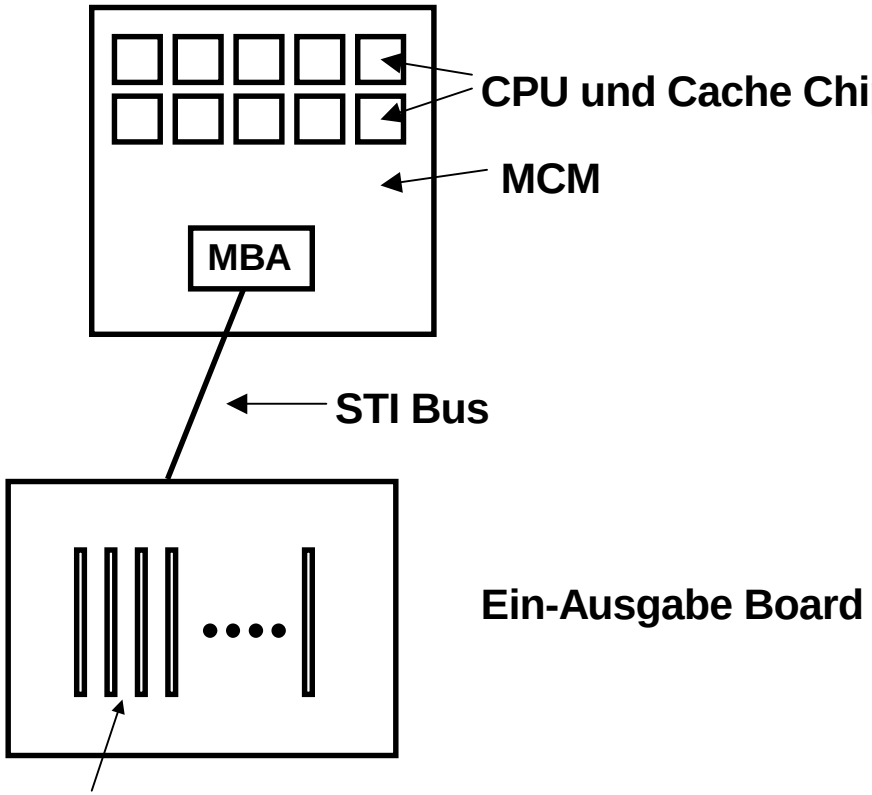

Ein-/Ausgabe Board

Auf dem Multichip Module sitzen neben den CPU- und Cache Chips vier Memory Bus Adapter (MBA) Chips, die eine ähnliche Rolle wie das Southbridge Chip in einem PC übernehmen. Aus jedem MBA werden sechs STI Busse herausgeführt, vergleichbar mit den PCI Bussen in einem PC. Die STI Busse münden in Ein-/Ausgabe Boards, die über STI Slots für die Aufnahme der Ein-/Ausgabeadapter Karten verfügen.

Die wichtigsten E/A Adapter Karten Typen haben Anschlüsse für

ESCON Kanal
FICON Kanal
OSA Adapter für Ethernet, ATM, ...


ALLIANCE

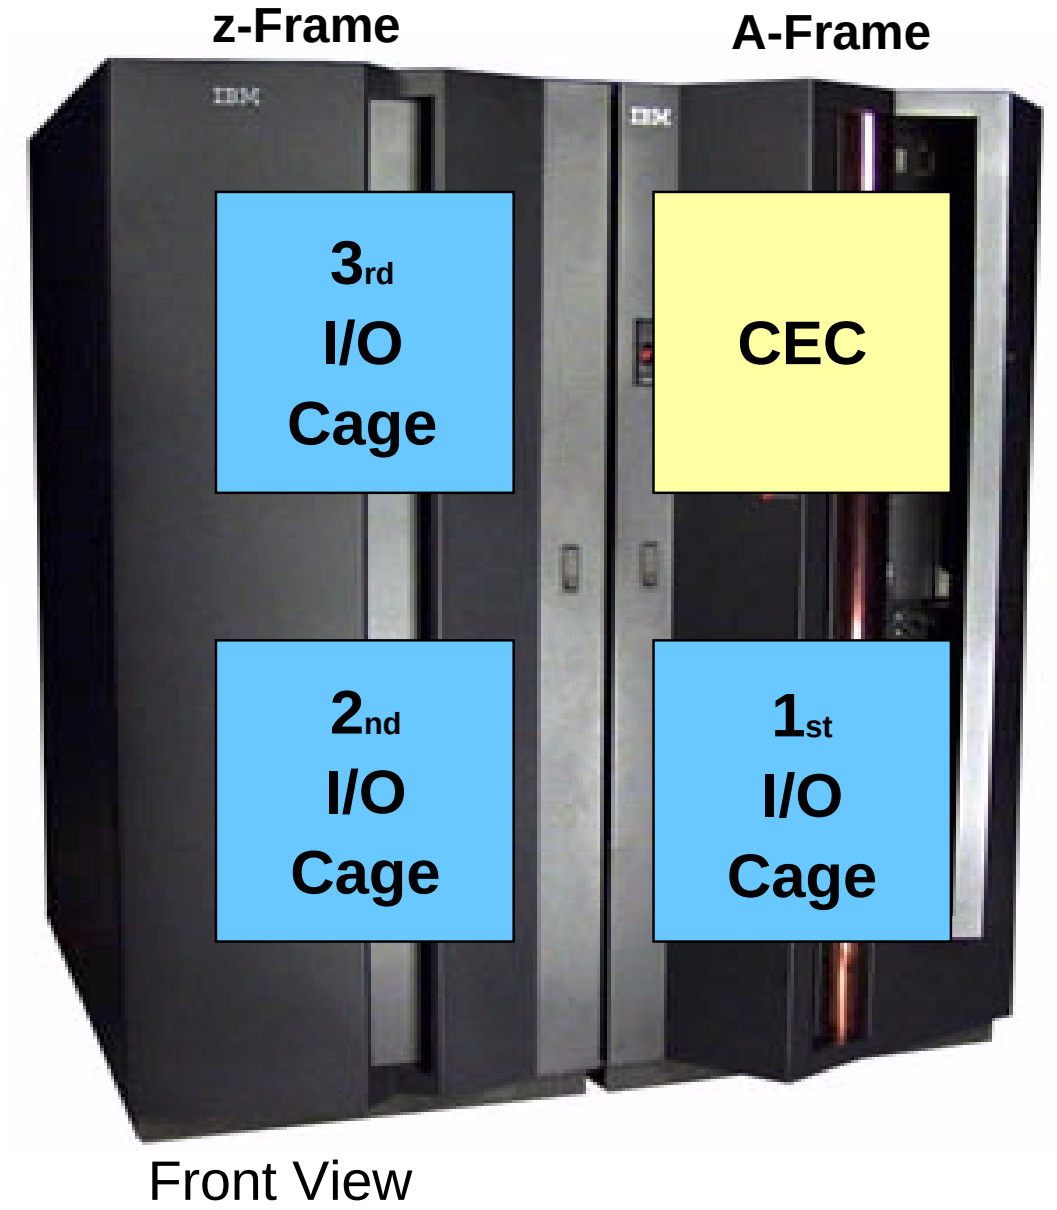

Front View

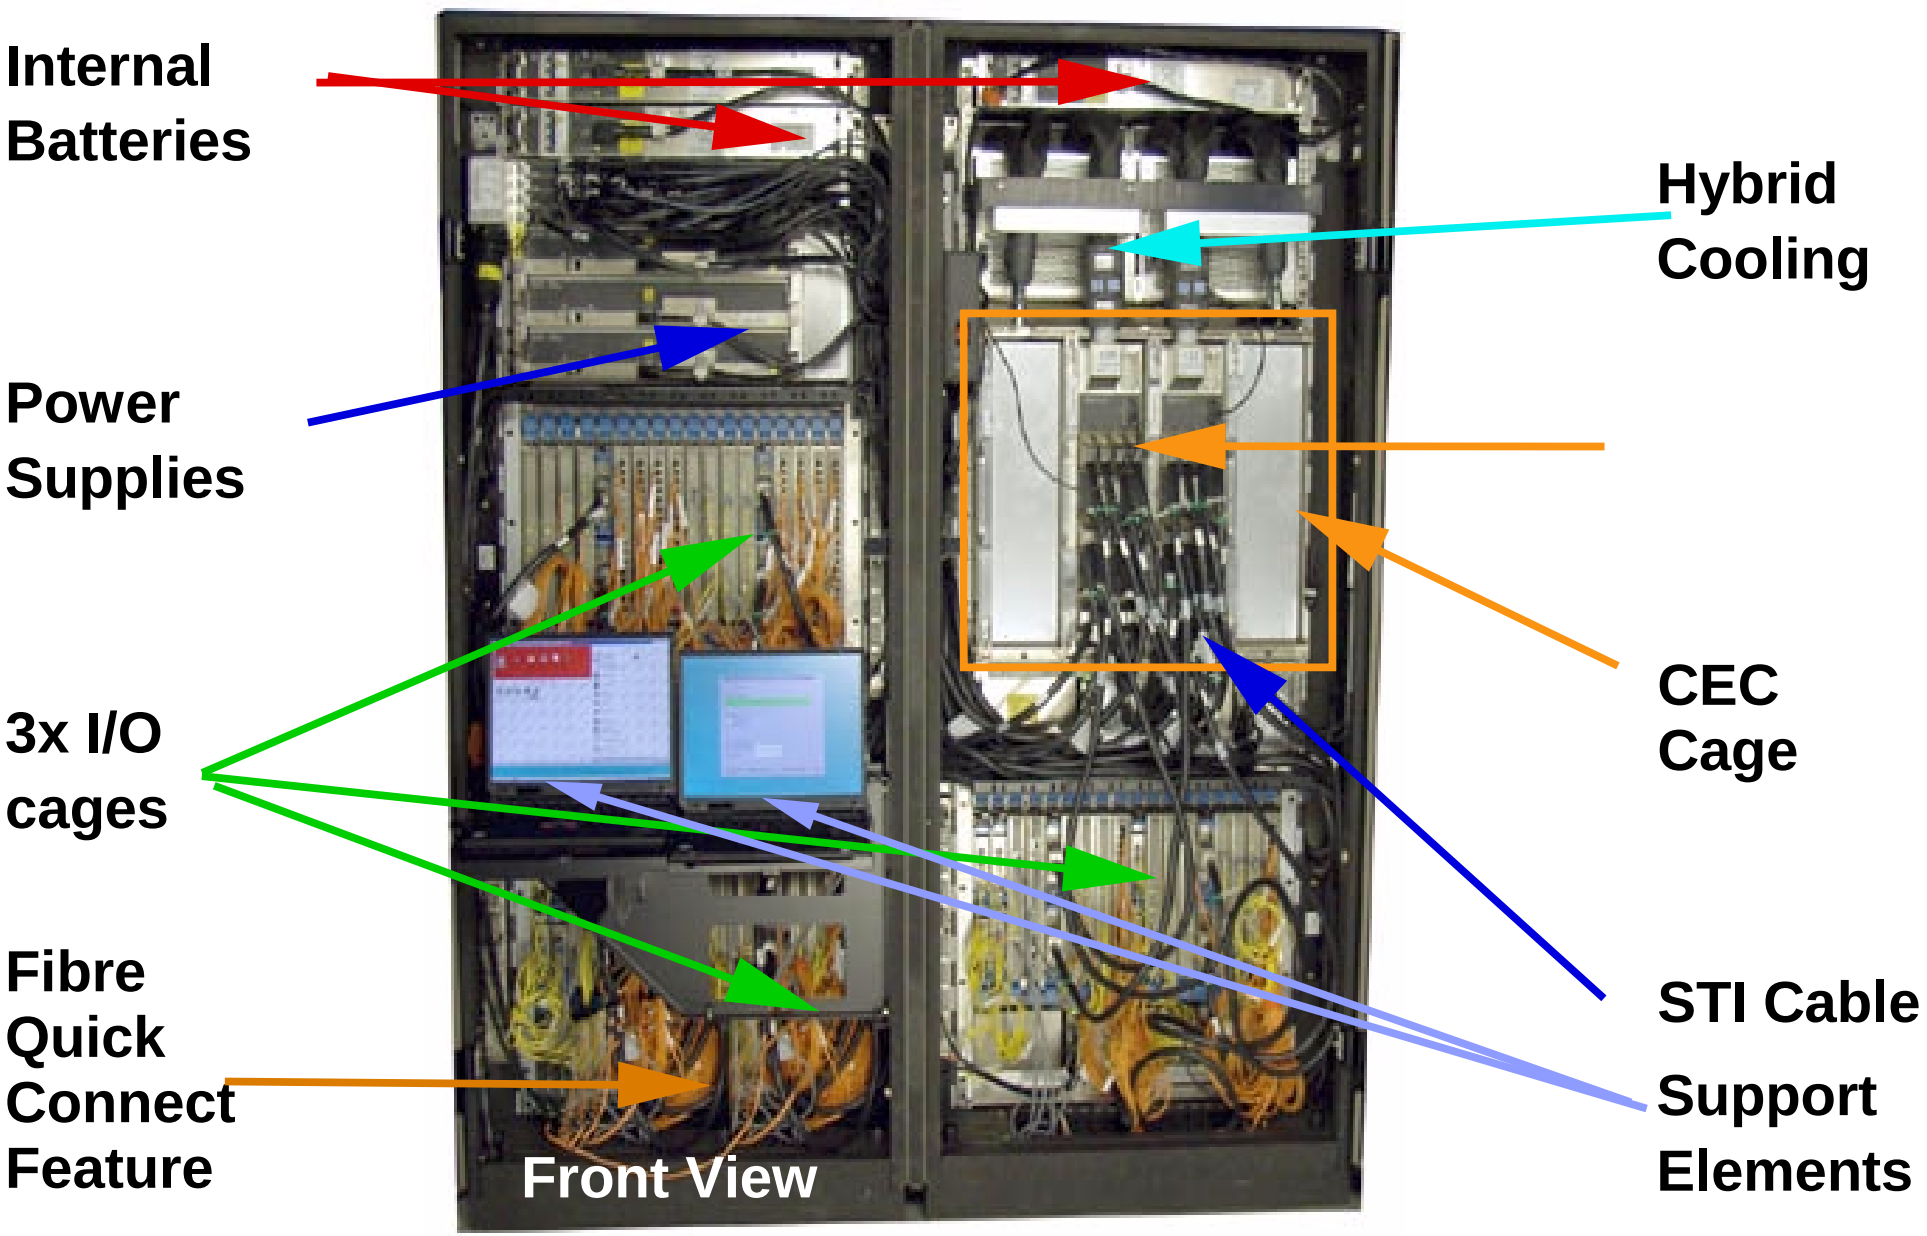
Internal Batteries
Power Supplies
3x I/O cages
Fibre Quick Connect Feature
Hybrid Cooling
CEC Cage
STI Cable
Support Elements
Front View


D46
D47
eserver

z9-109 'Central Electronic Complex': PUs, Speicher, I/O-Anschlüsse...

- Nimmt bis zu vier „Processor-Books“ auf

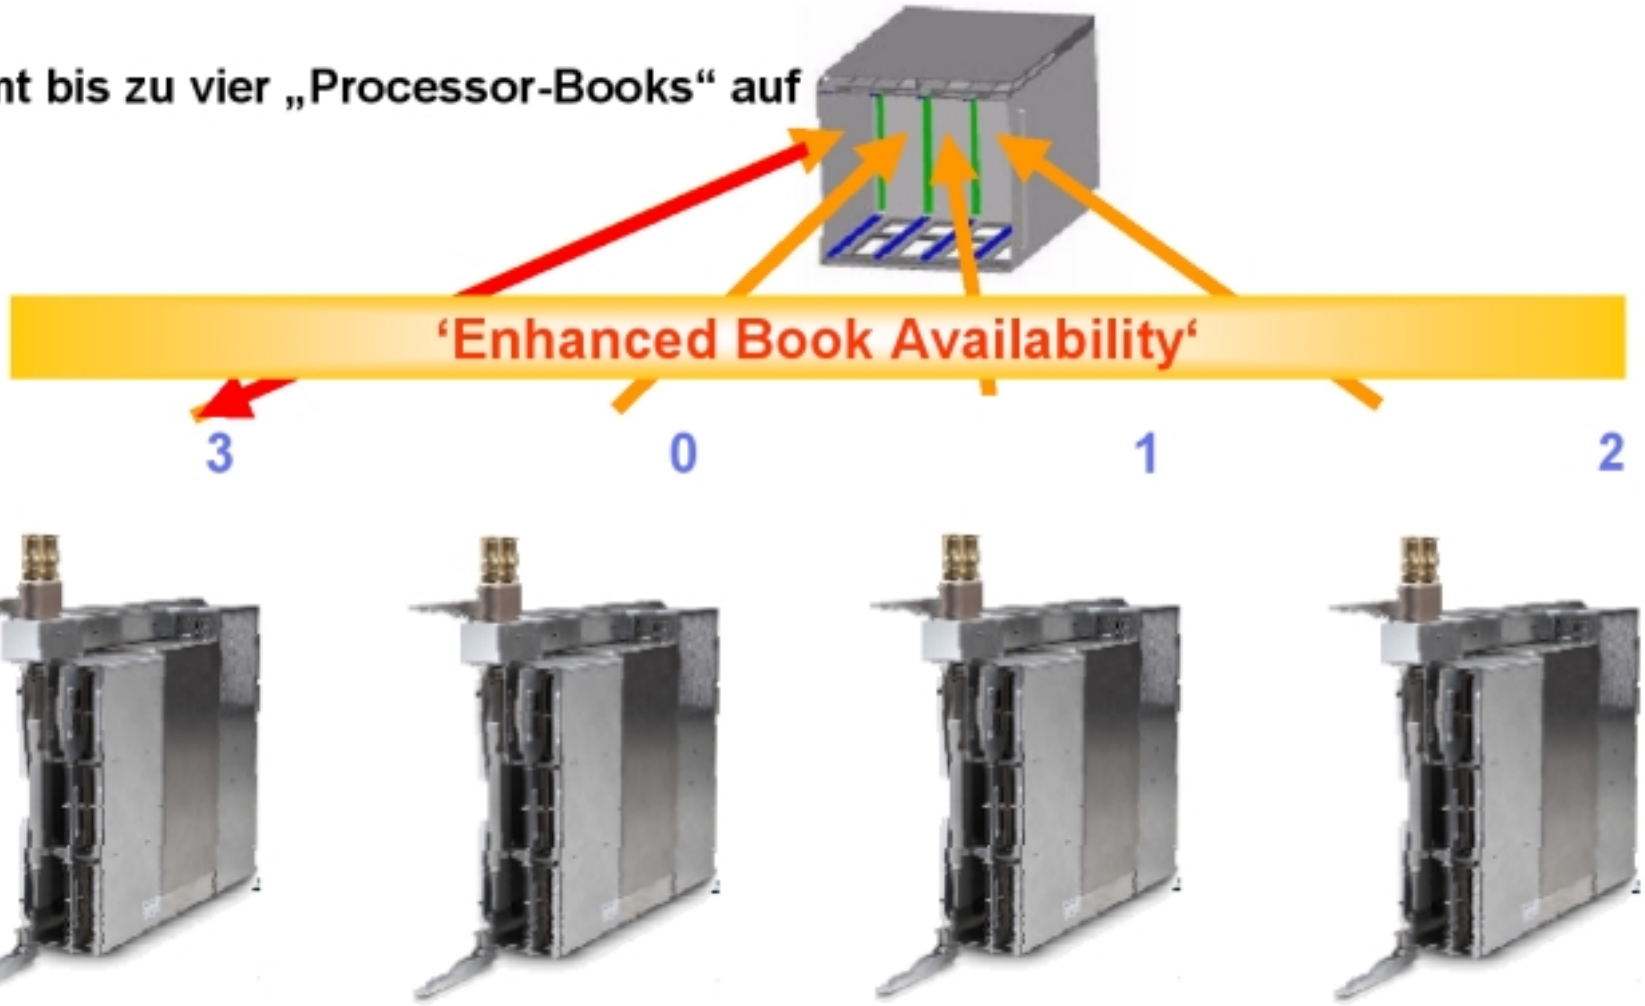

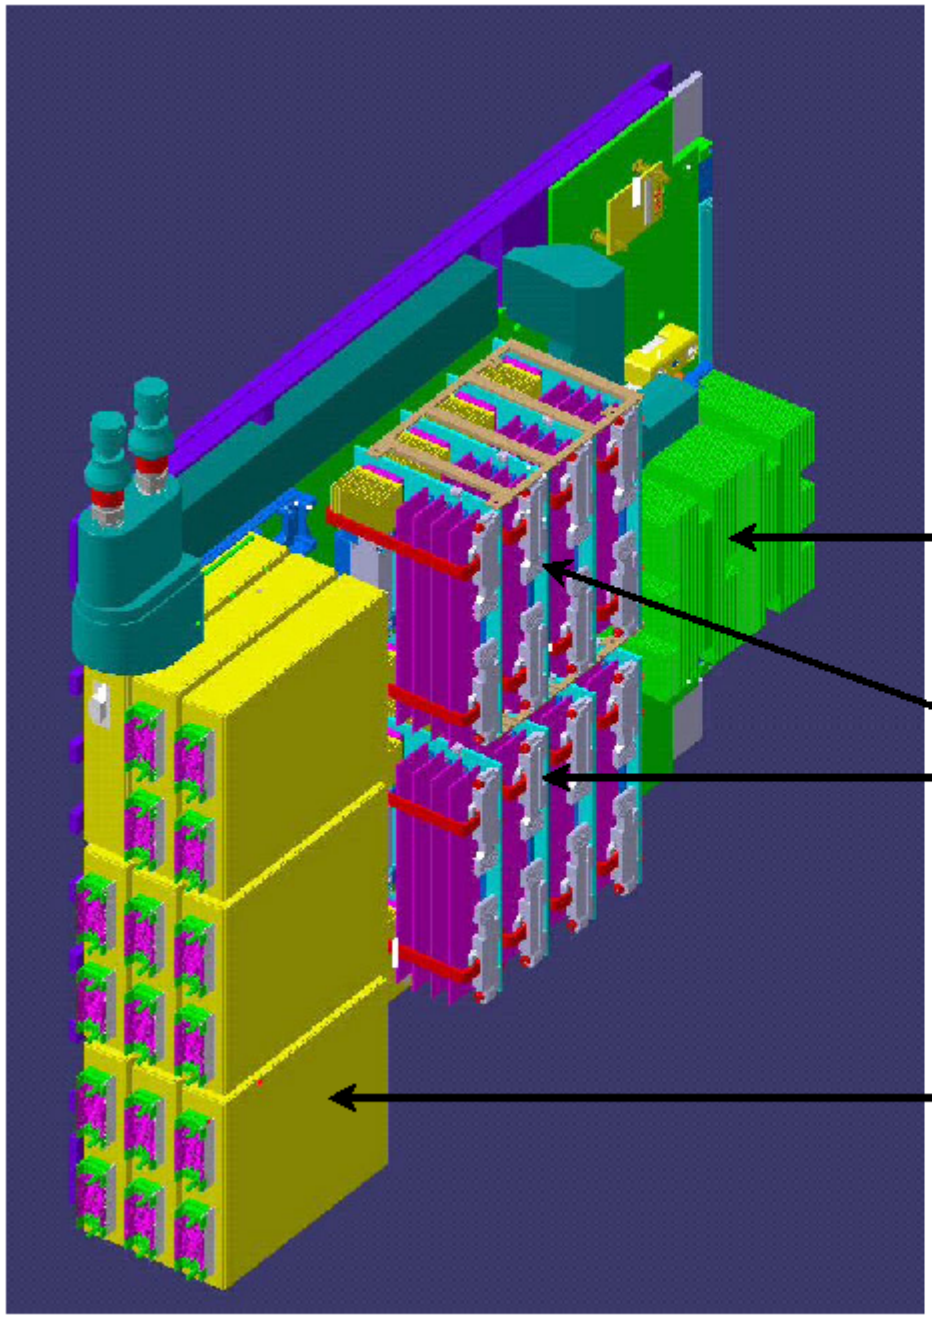

There are up to 16 STI buses per book to transfer data and each STI has a bidirectional bandwidth of 2.7 GByte per second. A four-book z9 EC server may have up to 64 STIs.

z9-109 Prozessor Book Layout

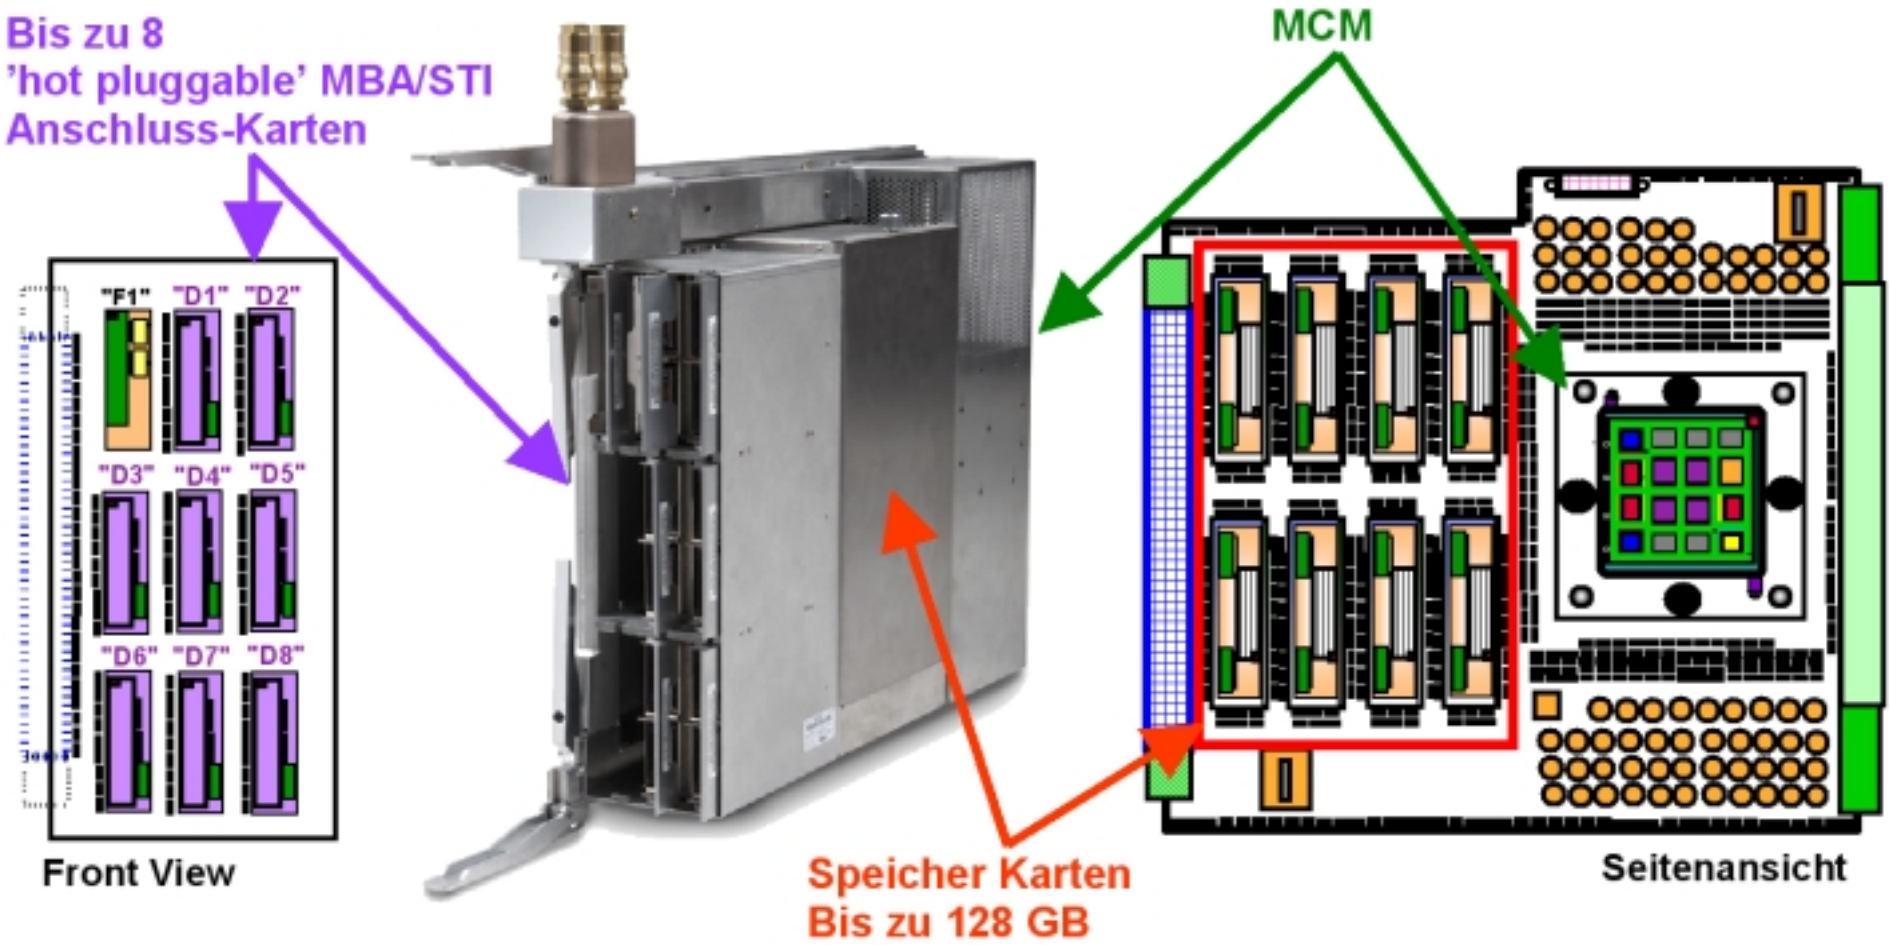

Each MBA fanout card connects to up to two STI cables. There are up to 8 MBA fanout cards per book, each driving two STIs, resulting in 16 STIs per book. All 16 STIs in a book have a data rate of 2.7 GB per second. A maximum of 1024 channels per server is supported.

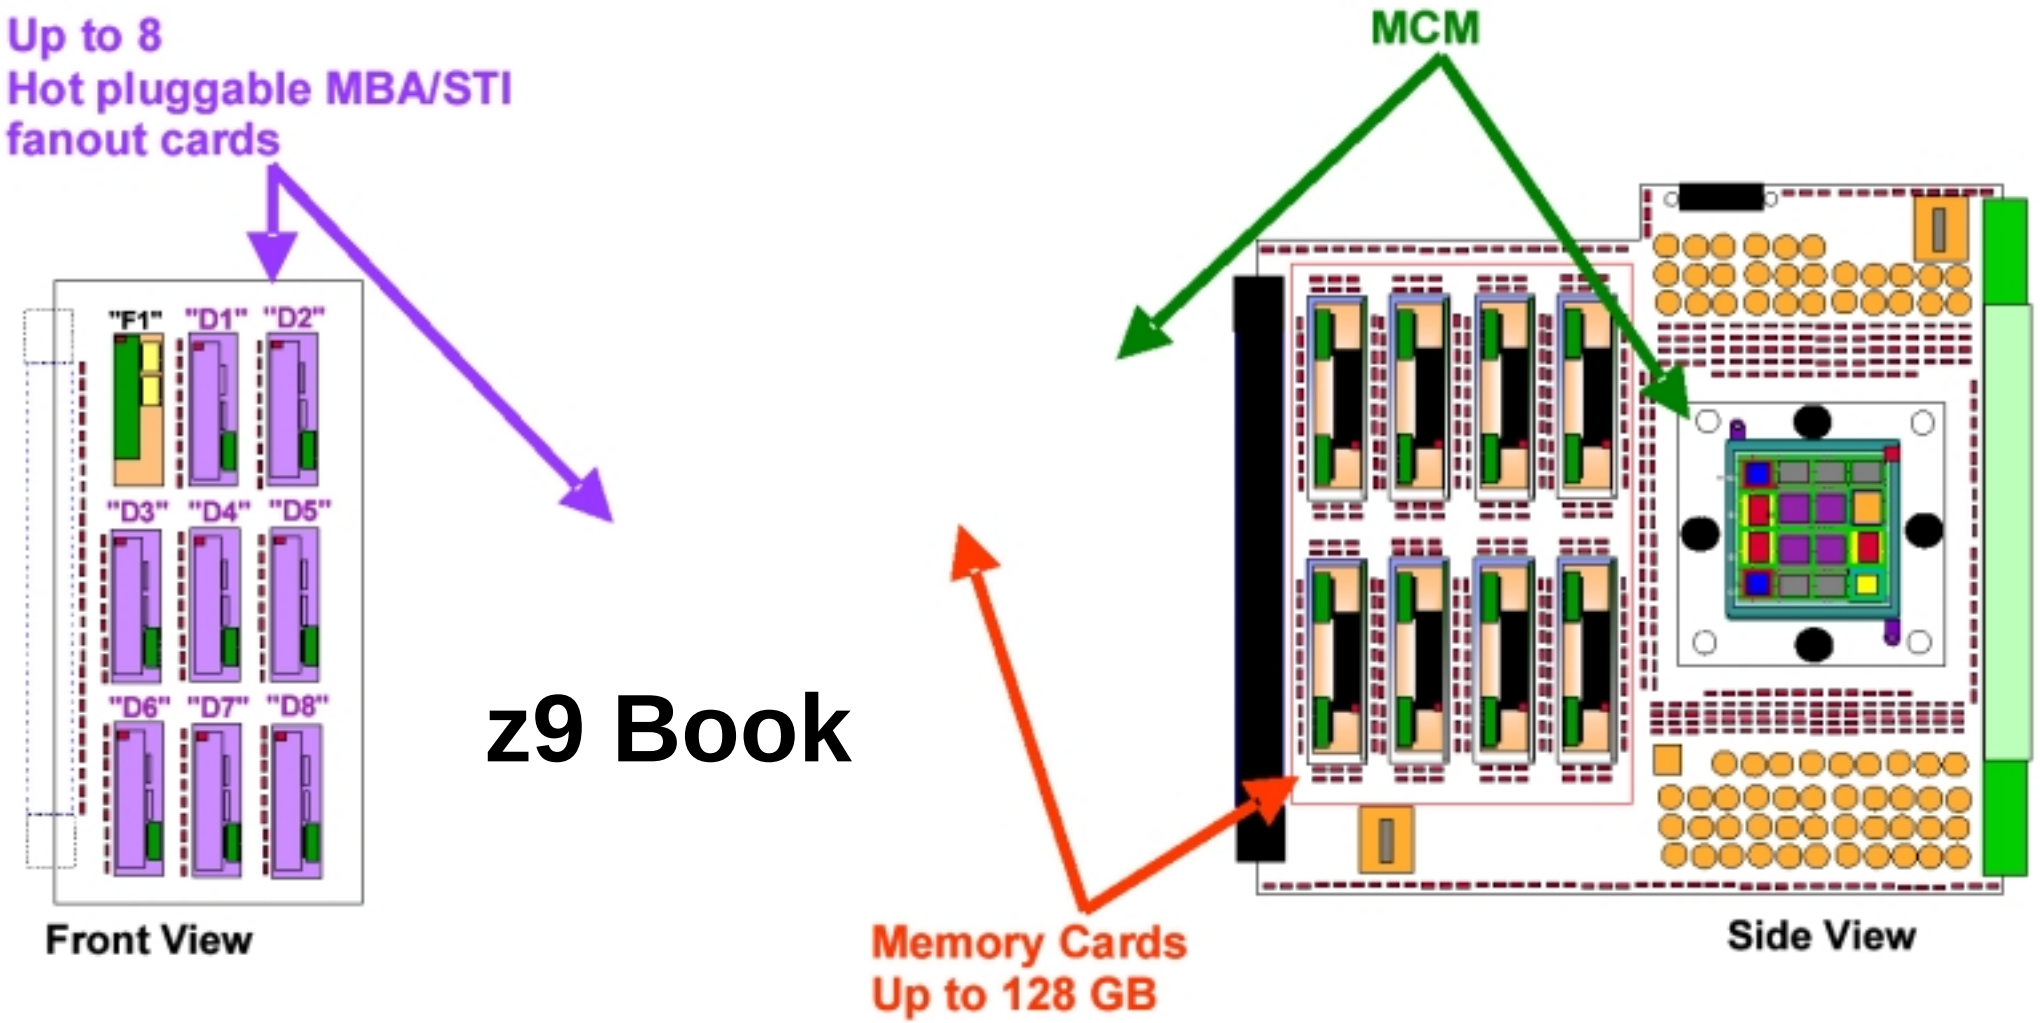

Connectivity from the CPU/memory area to the I/O area is established via eight MBA/STI hot pluggable cards at the front of the book. Each of the eight MBA/STI cards provides two 2.7 GByte/sec bidirectional interfaces to the I/O cage realized by the big black cables, i.e.,the aggregate external bandwidth is more than 43 GB/sec.

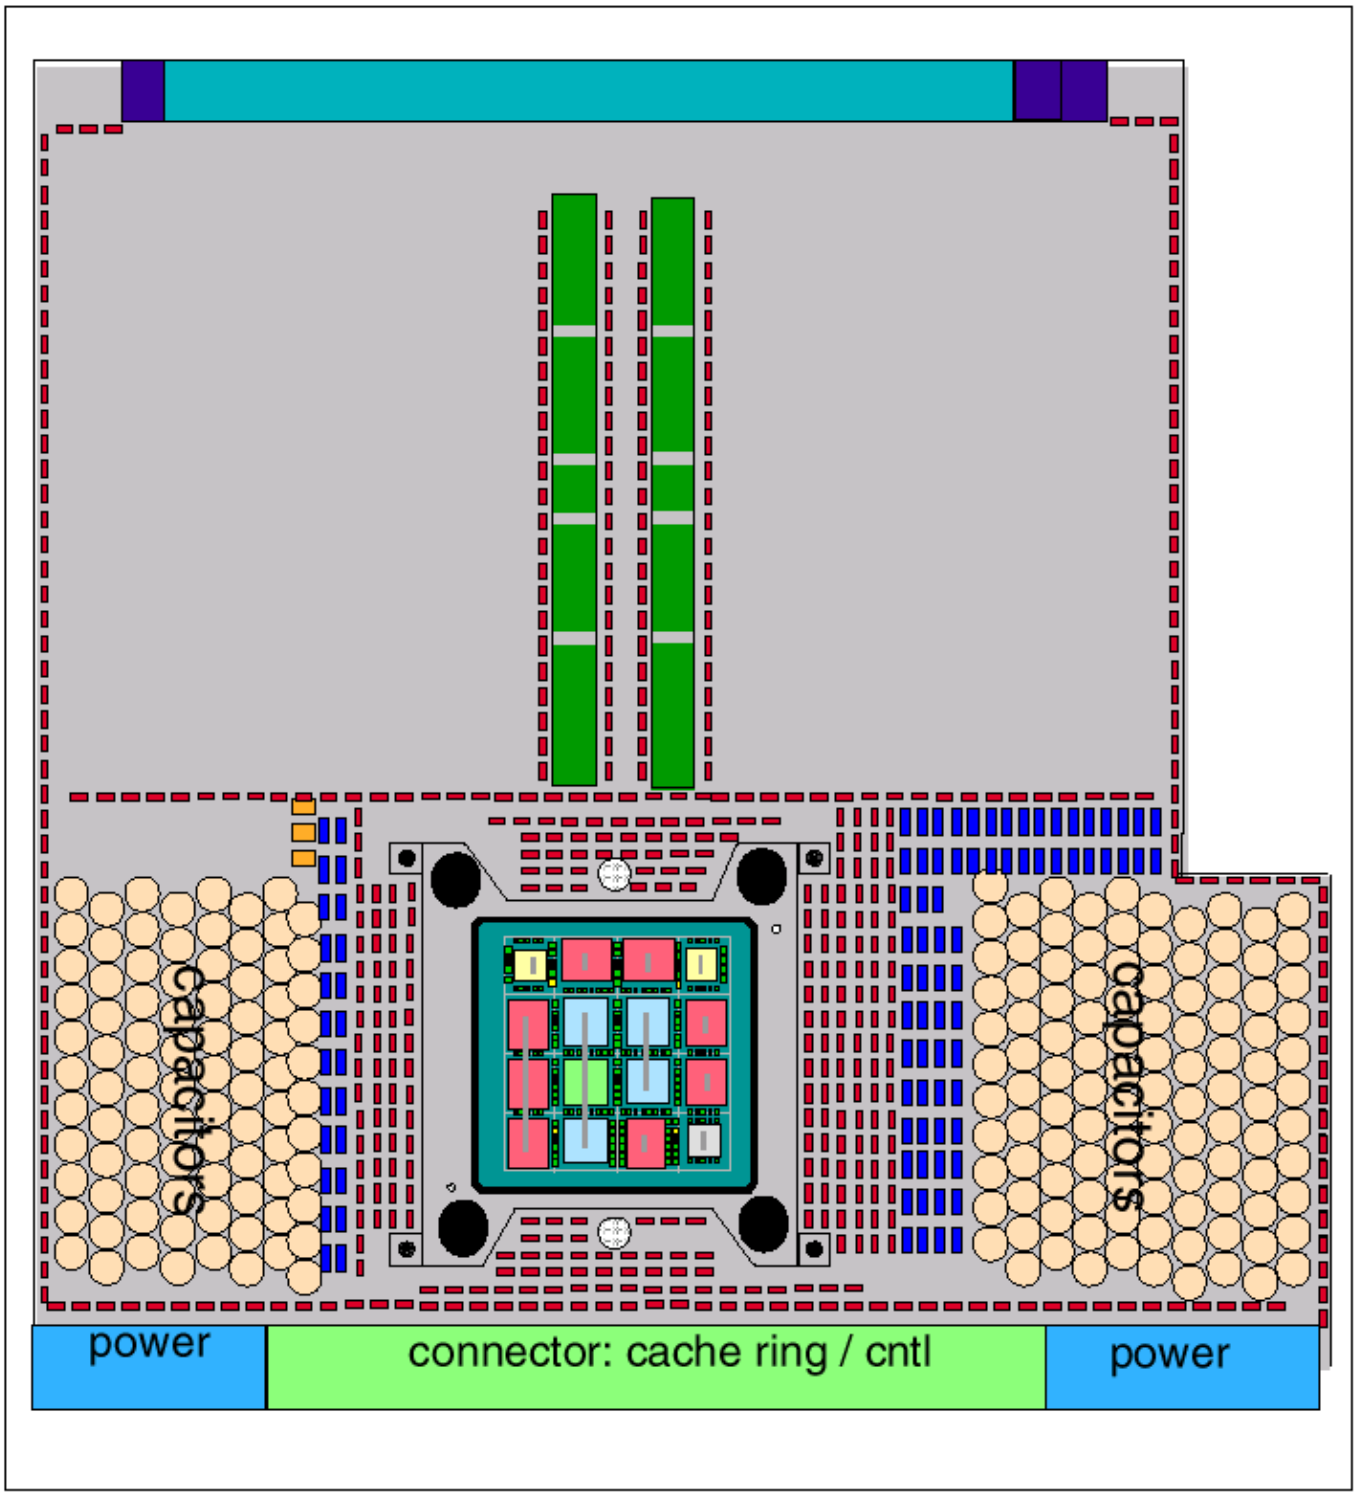
capacitors
capacitors
power
connector: cache ring / cntl
power

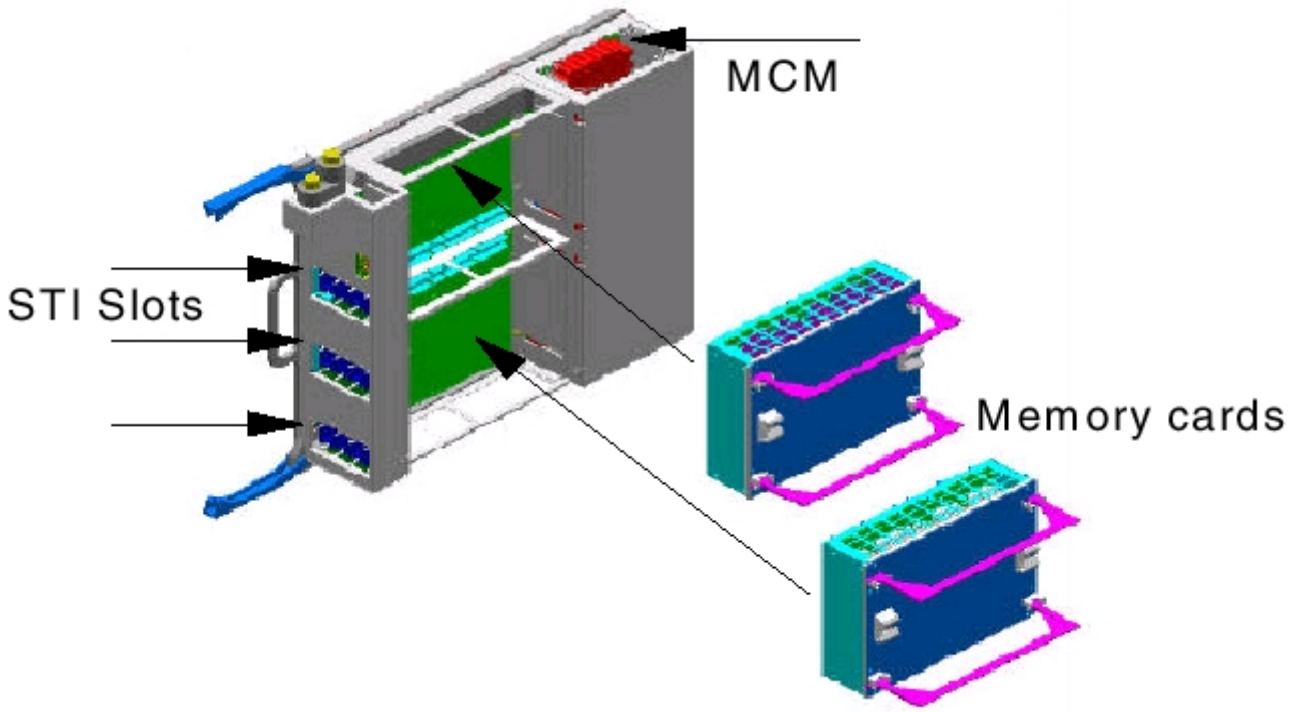

z9 und z990 Processor Book

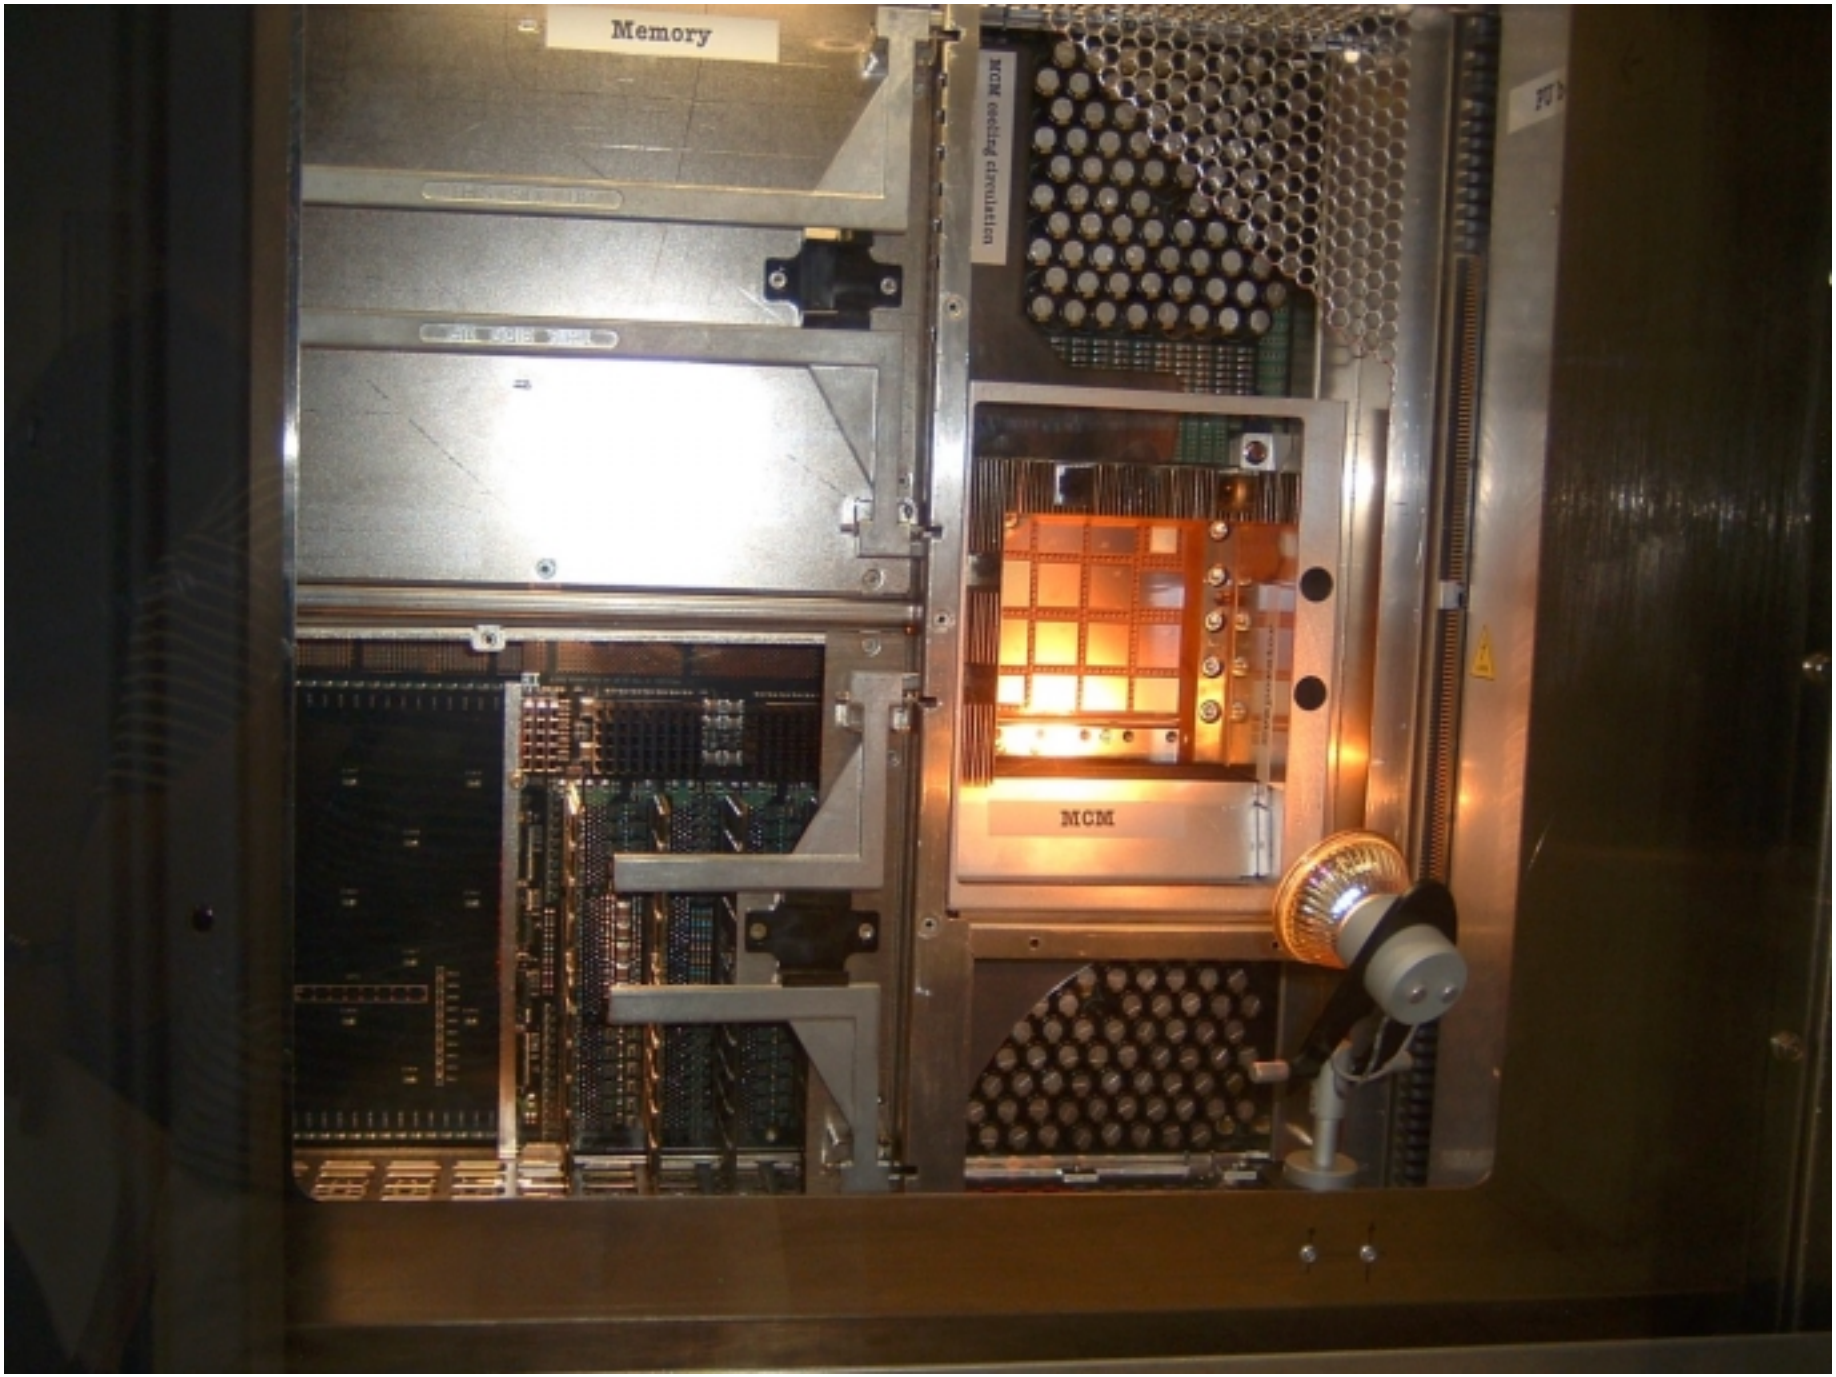
Memory
MCM cooling circulation
MCM

MCM

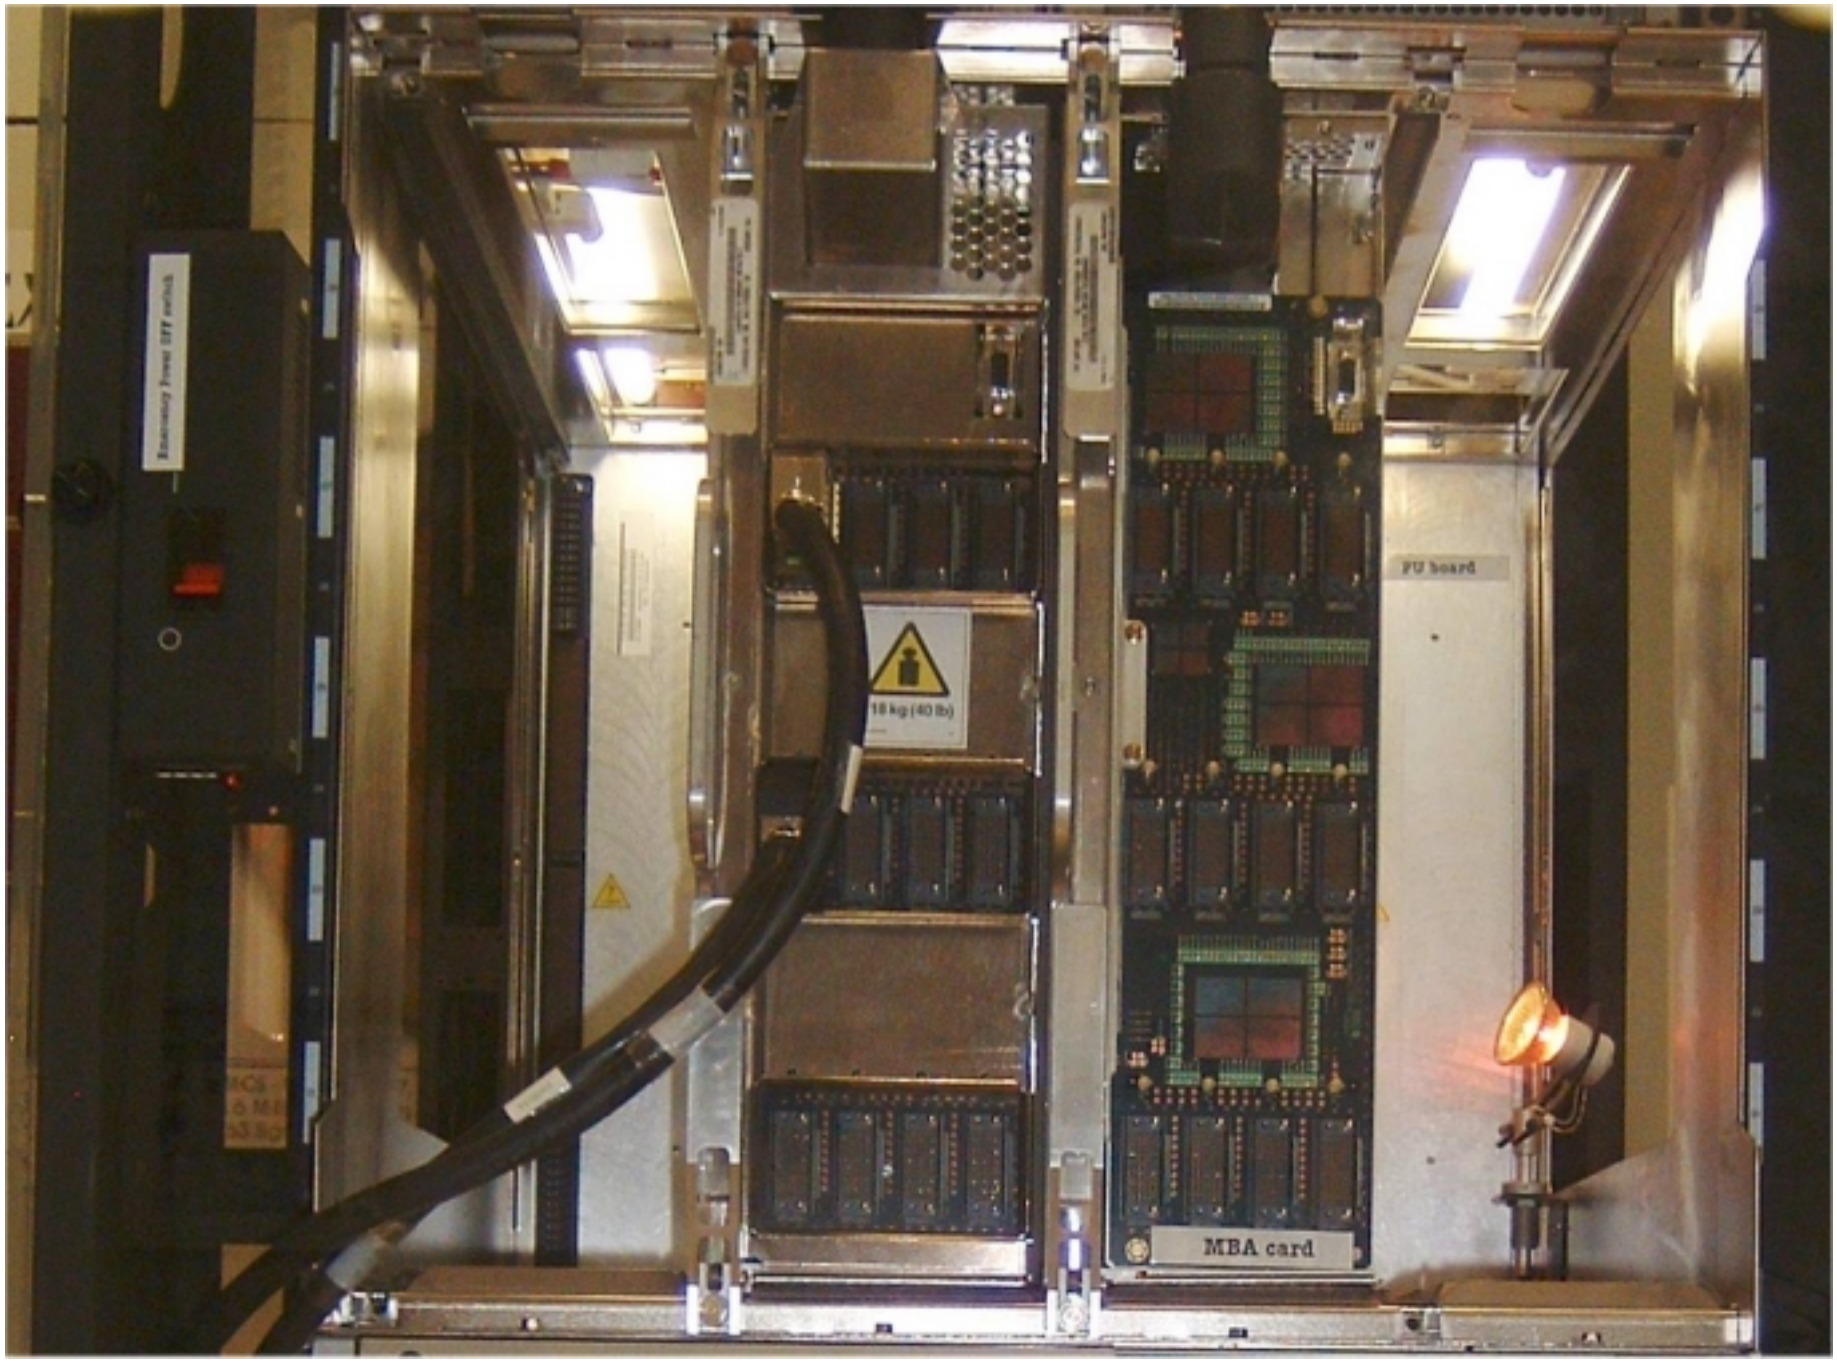
PU board
MBA card

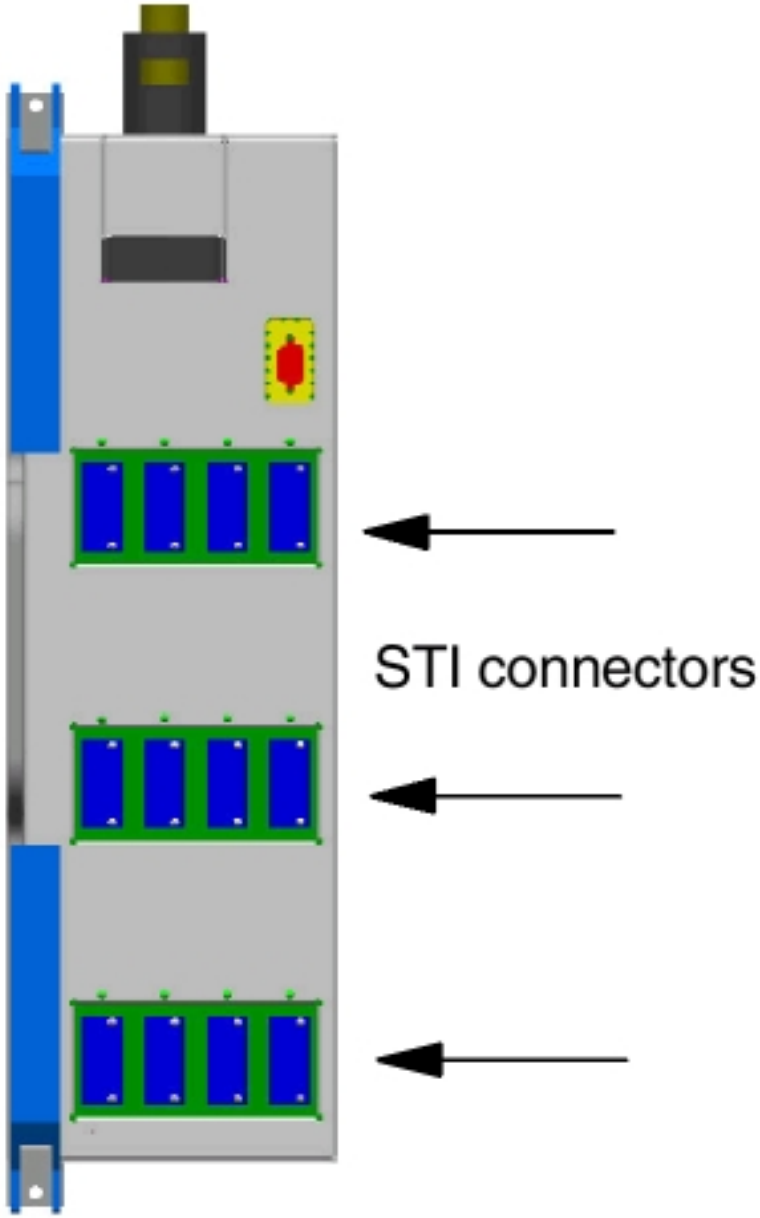

Book front view

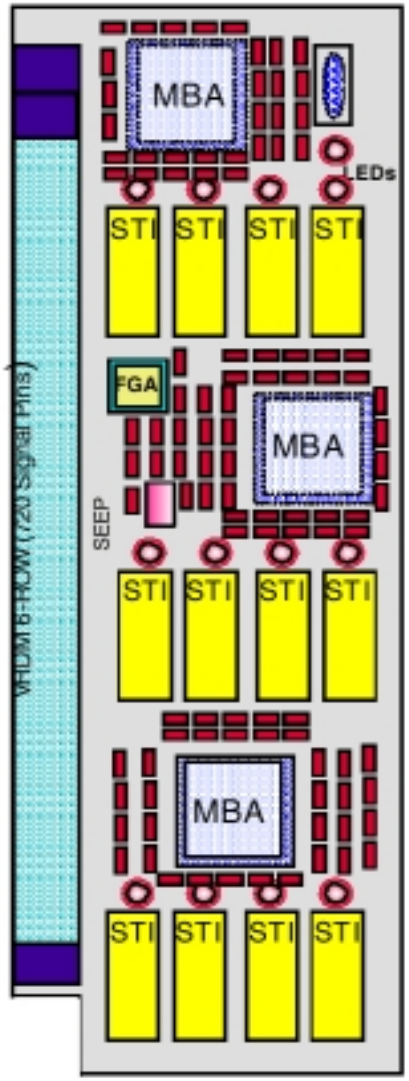

MBA card

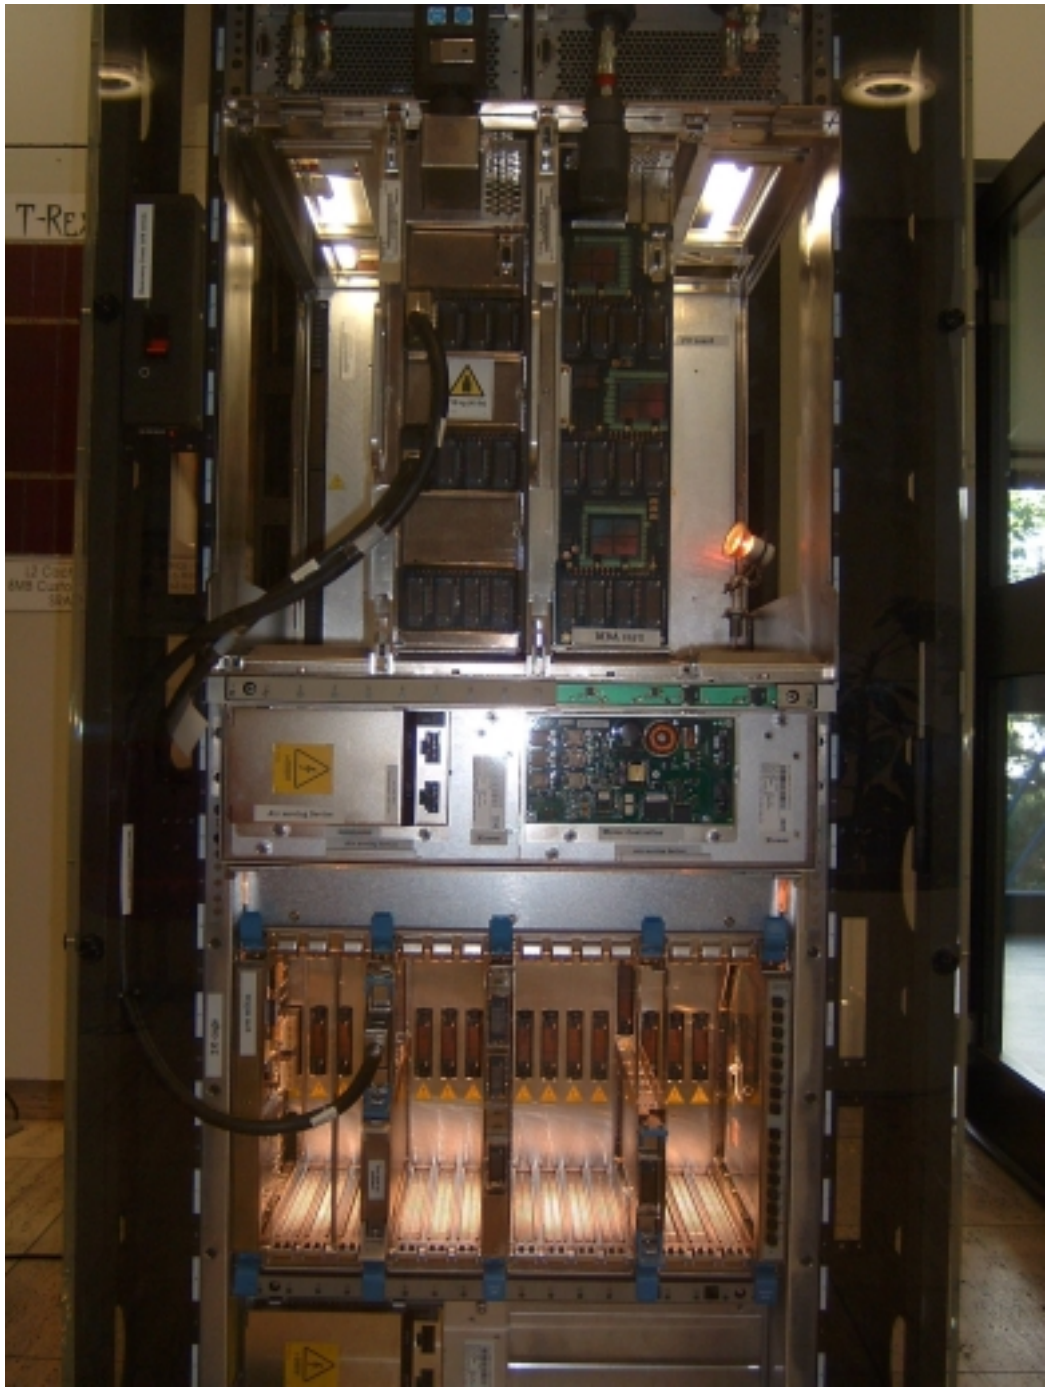

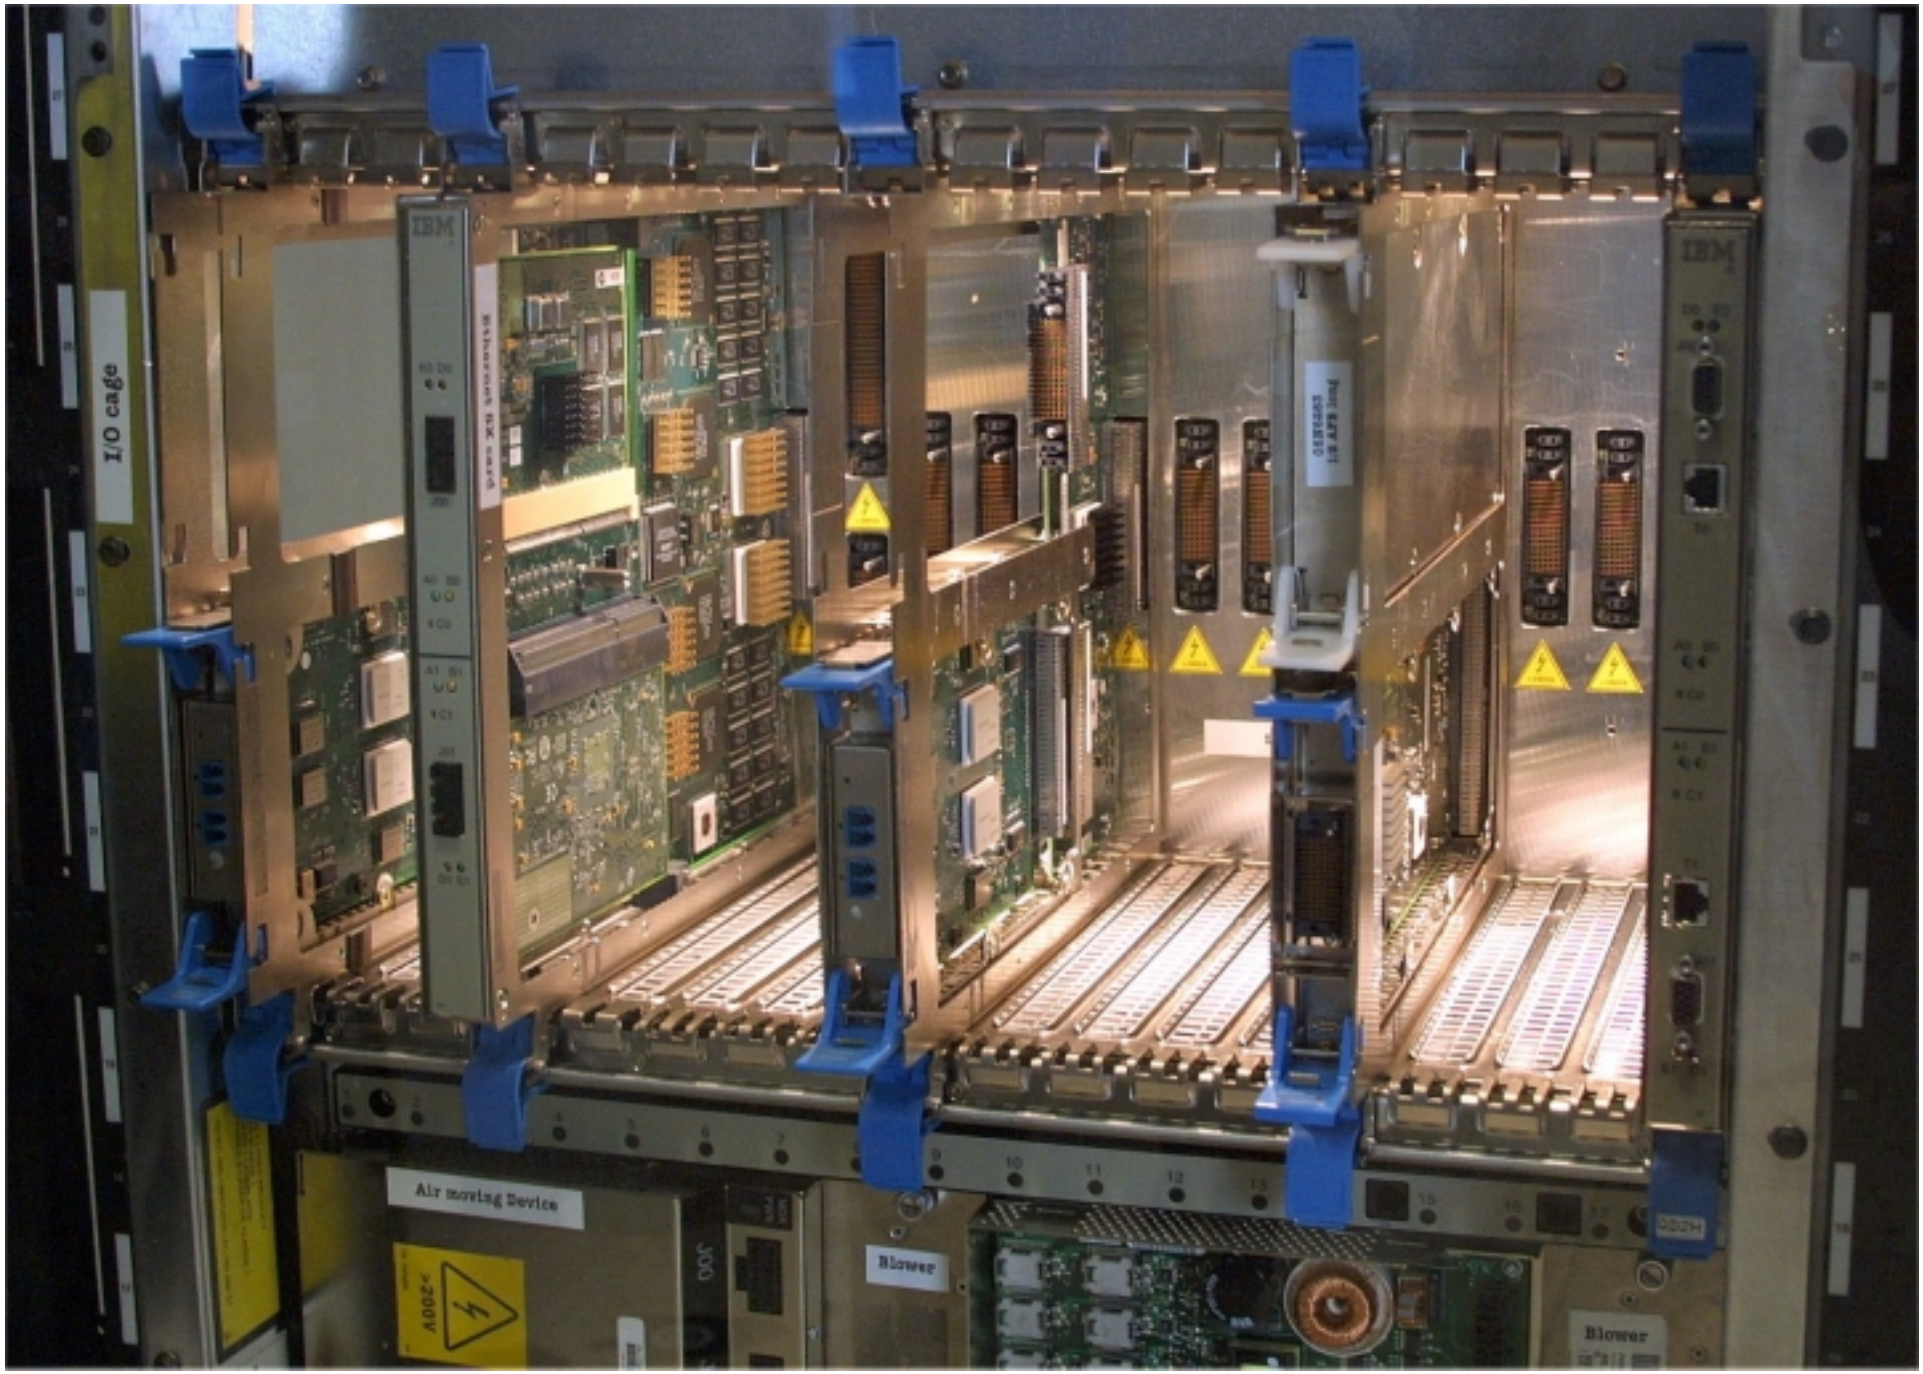
I/O cage
IBM
Ethernet SX card
Air moving Device
Blower
Blower

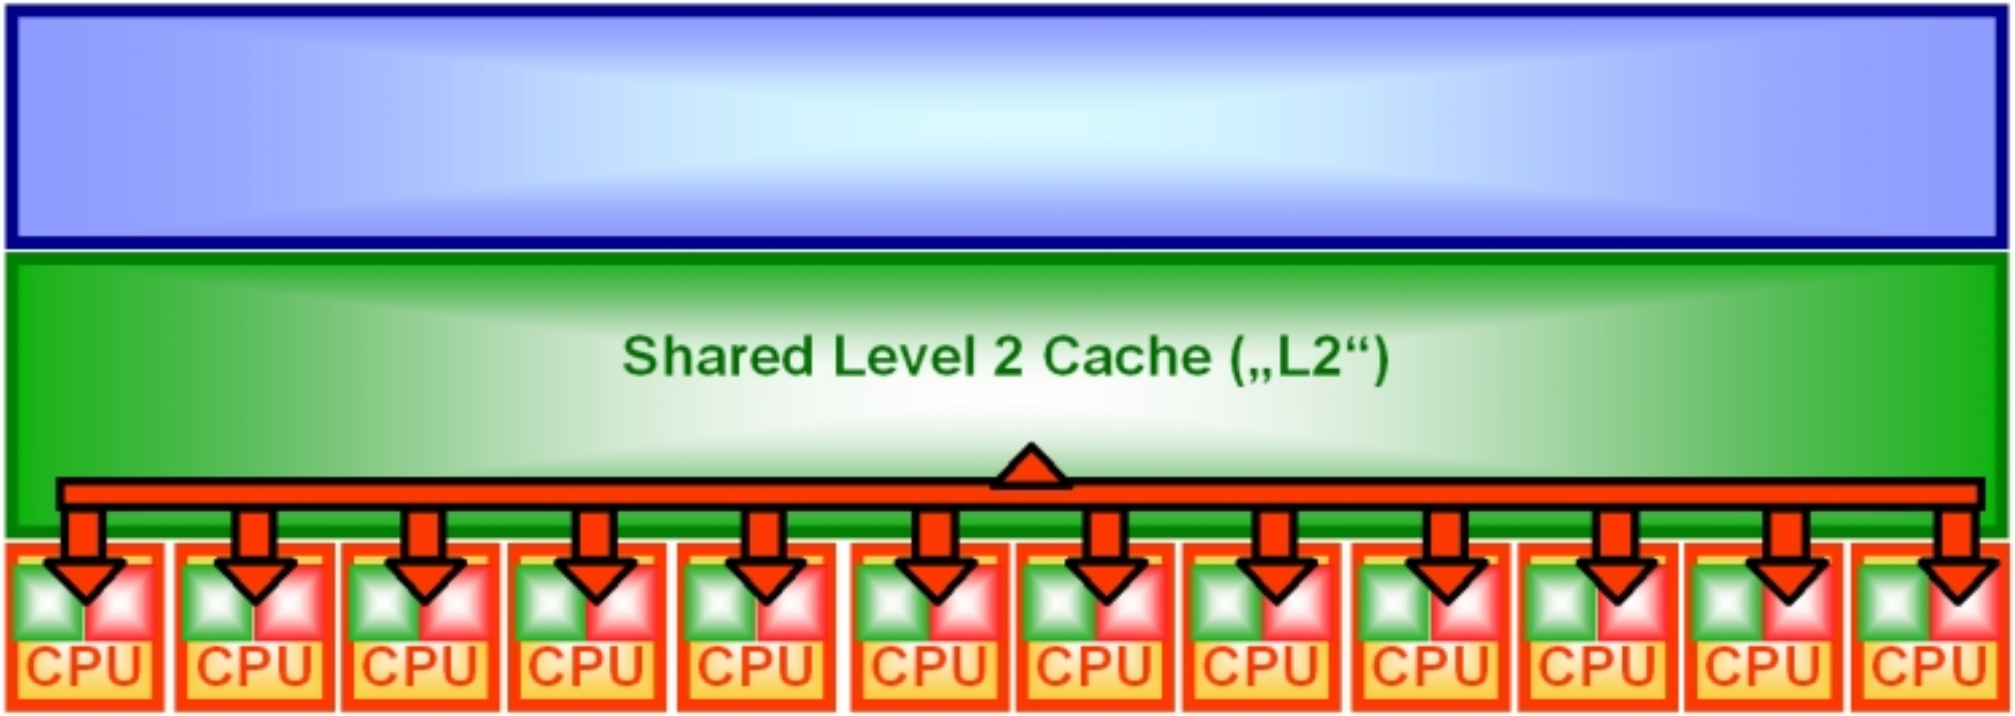

zSeries - Zentraler Switch
mit gleichzeitigem Zugriff aller Prozessoren

z990 Model D32 (GA2)

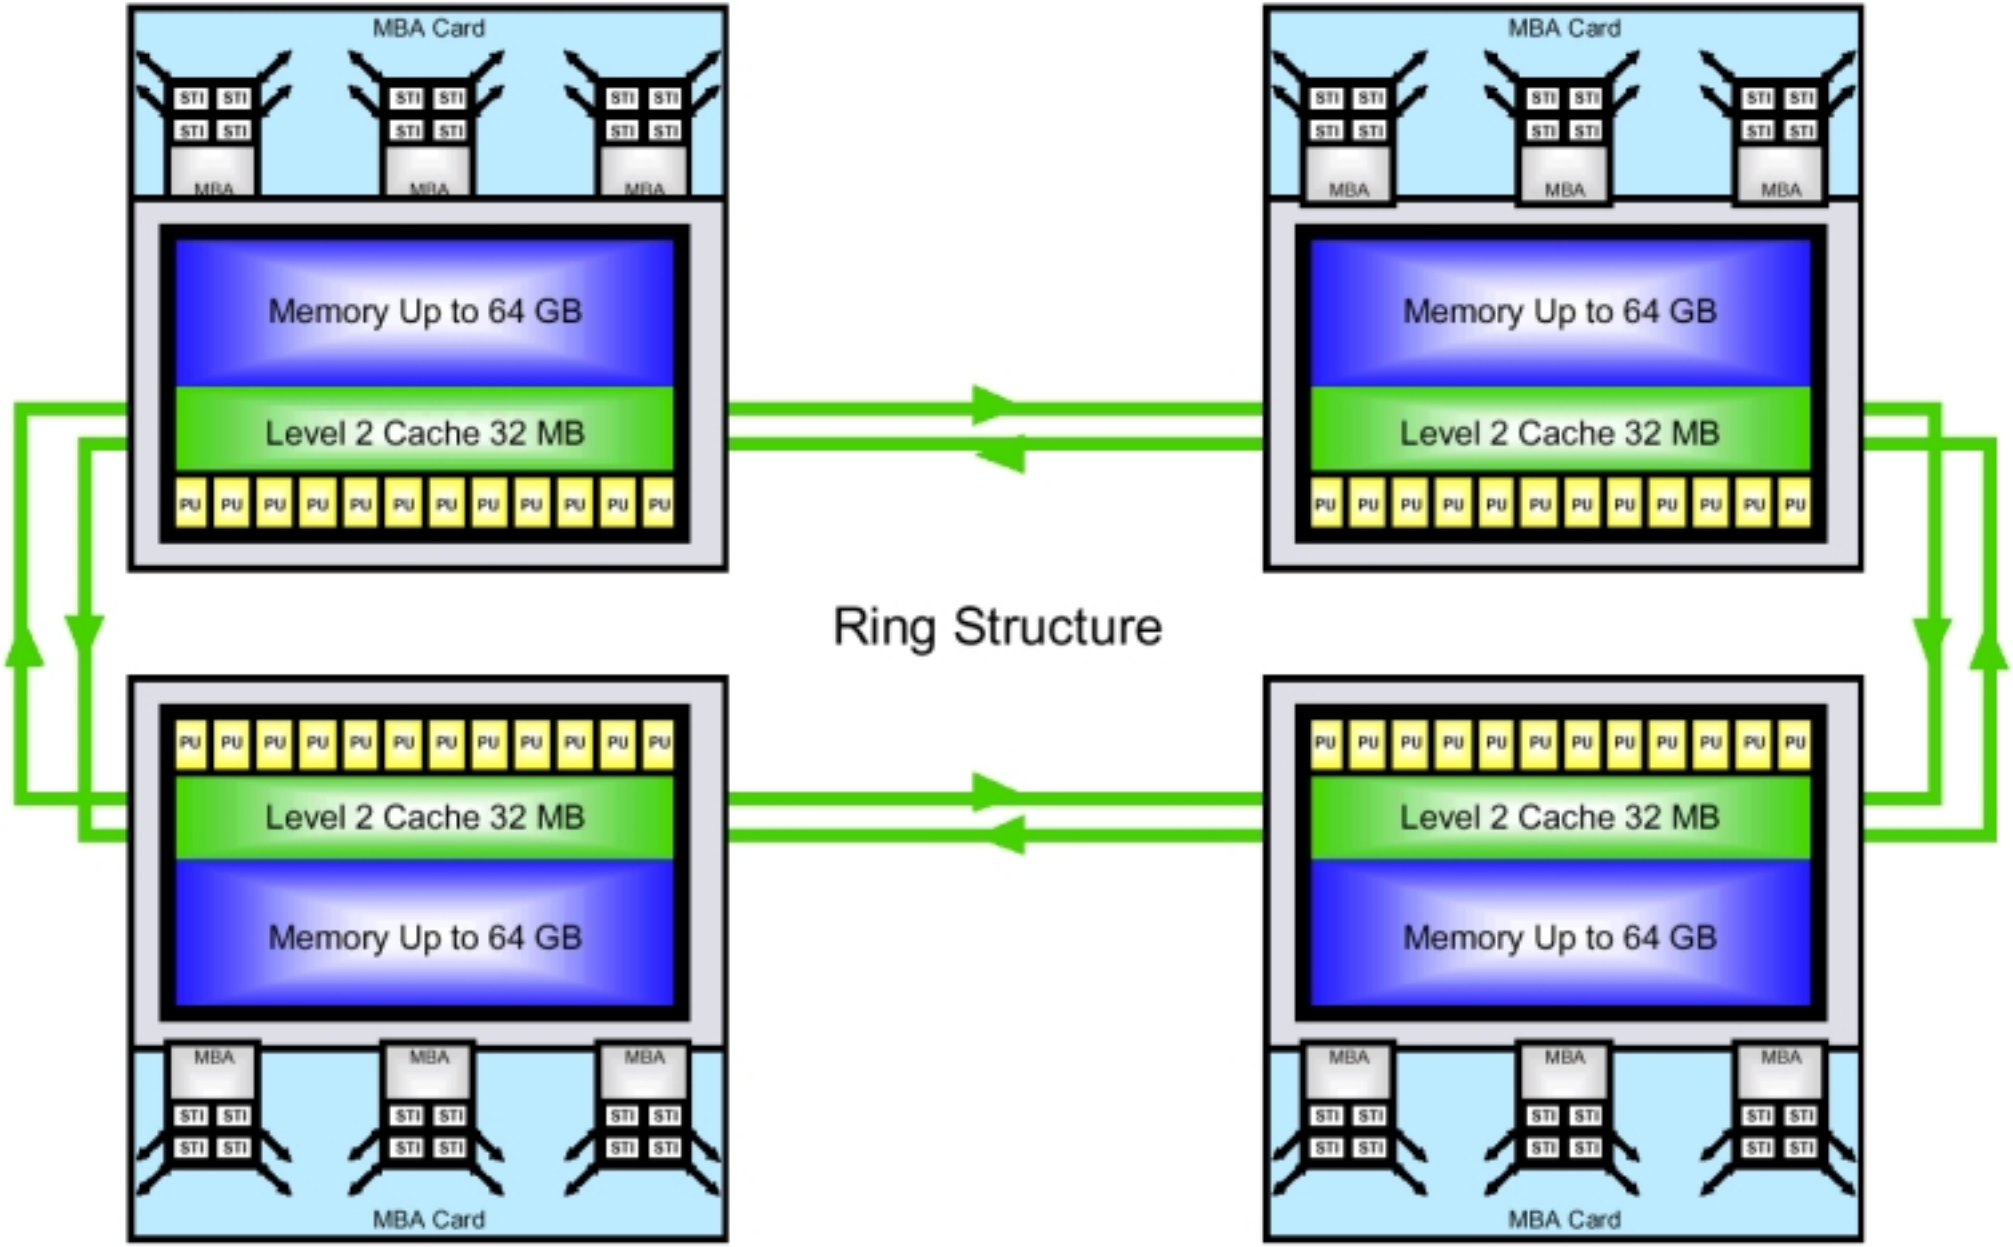

z9-109 Prozessor-Book: Logische Systemstruktur

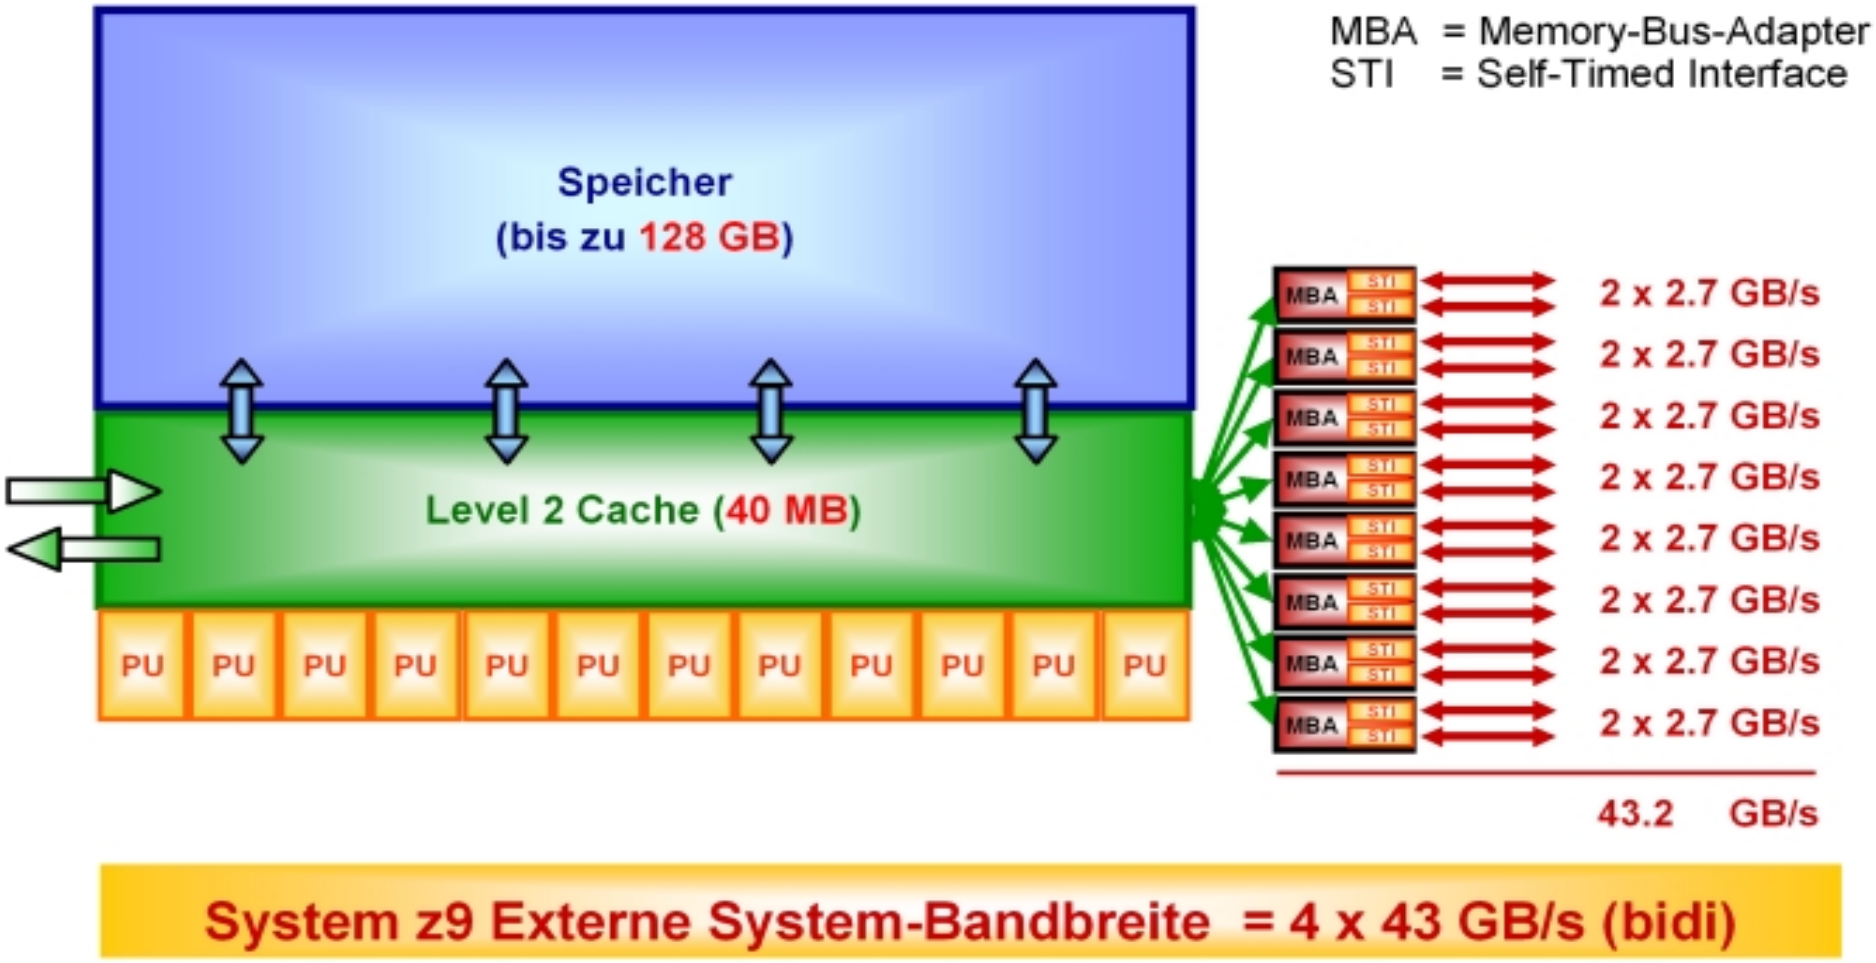

Im Gegensatz zu allen anderen Architekturen kommunizieren die E/A Geräte mit dem L2 Cache und nicht mit dem Hauptspeicher. Die zSeries Ingenieure waren in der Lage, die resultierenden Cache Koheränz Probleme zu lösen.