GRAPE Project: special-purpose computers for many-body simulations

Jun Makino

University of Tokyo

Talk Overview

1. Why to develop special-purpose computers?
   - Technology trends of general-purpose computers
   - Potential advantage of SPCs
2. The GRAPE approach
3. Past and Present of GRAPE systems
4. GRAPE-6: The next generation

Technology trends of GPCs

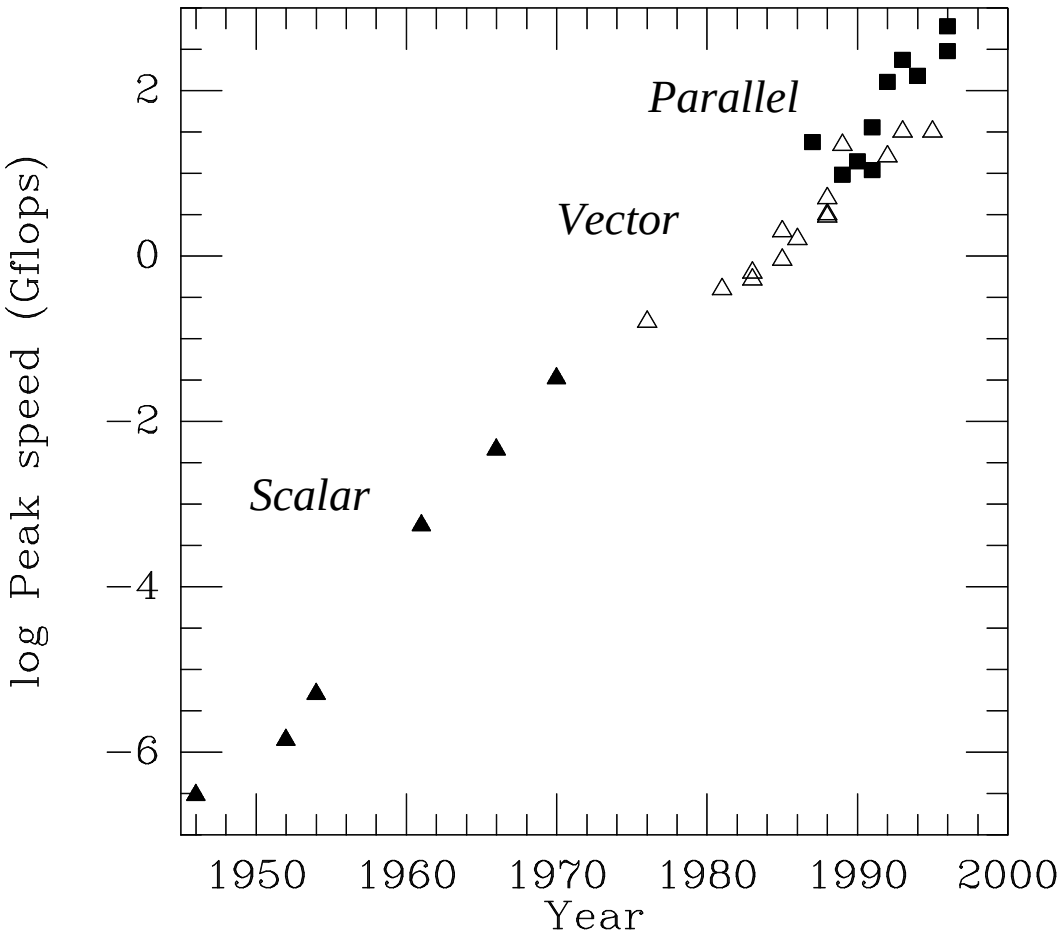

- 100 x faster in 10 years
- Evolution of Archtecture (scalar→vector →parallel)

Technology trends of GPCs (continued)

Hardware efficiency goes down *exponentially*

- Available number of transisters : 100 $\times$ in 10 years
- Vector-parallel machines in the last 20 years: number of ALUs $\sim 300\times$, total transistor count $> 10^4\times$
- Microprocessors: very small increase in the number of operations per cycle in the last 10 years

Microprocessor Evolution

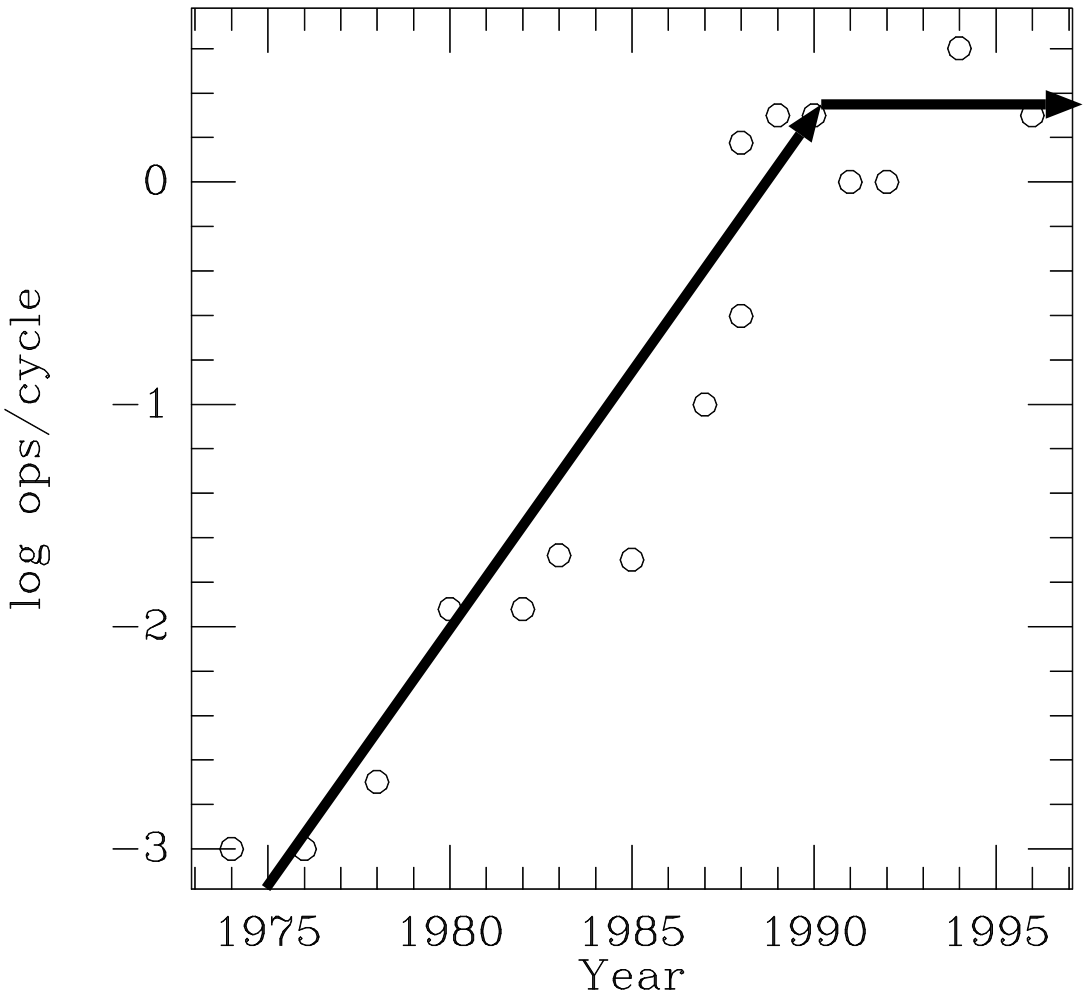

Number of floating point operations per cycle of representative microprocessors

Increase halted by 1990

Why efficiency goes down?

- memory requirement $\propto$ speed
  $\rightarrow$a fast processor with tiny memory is useless.
- It is difficult to improve the speed: "Memory Wall"

$\downarrow$

The processor-memory bandwidth determines the actual performance

Potential advantage of SPCs

Can achieve better C/P by eliminating hardware overhead

IF

- large memory is not necessary
- many arithmetic units can be used without requiring high memory bandwidth

Practical advantage in C/P

Most ideal case:

- relative advantage in 1998:100–1000
- future: relative advantage increases by a factor of 10 per decade

Practical difficulties of SPCs

- Large amount of resources (time, budget) is necessary for initial development: Potential advantage can be easily lost.
- (Because it is SPC), there will be limitations in problem range and available algorithms.
- People who build SPCs are not experts...

Past/present examples of SPCs

- Ising machines
- QCD machines
- Machines for classical particle-based simulations

Particle-based simulations

Master equation:

$$\frac{d^2 x_i}{dt^2} = \sum_{j \neq i} f_{ij}$$

- In many cases, interaction is central force (potential depends on distance only)
- Simulation timescale tends to be very long : small memory requirement
- Calculation of interaction is dominant : small bandwidth requirement

Machines for particle-based simulations

- Programmable parallel computers
  - Orrery, ATOMS, MUSIC,
  - GROMACS, HaMM
- Pipelines for interaction evaluation
  - DMDP, FASTRUN

Programmable parallel machines for particle-based simulations

—Pros

- Minimal requirement for memory, memory bandwidth, communucation bandwidth etc.

—Cons

- Time necessary for deveopment: machines may become obsolete before completion.
- No OS, no compilers: Software development also takes time.

Pipelined processor for interaction evaluation

DMDP (A. Bakker *et al.*, Delft Institute of Technology)

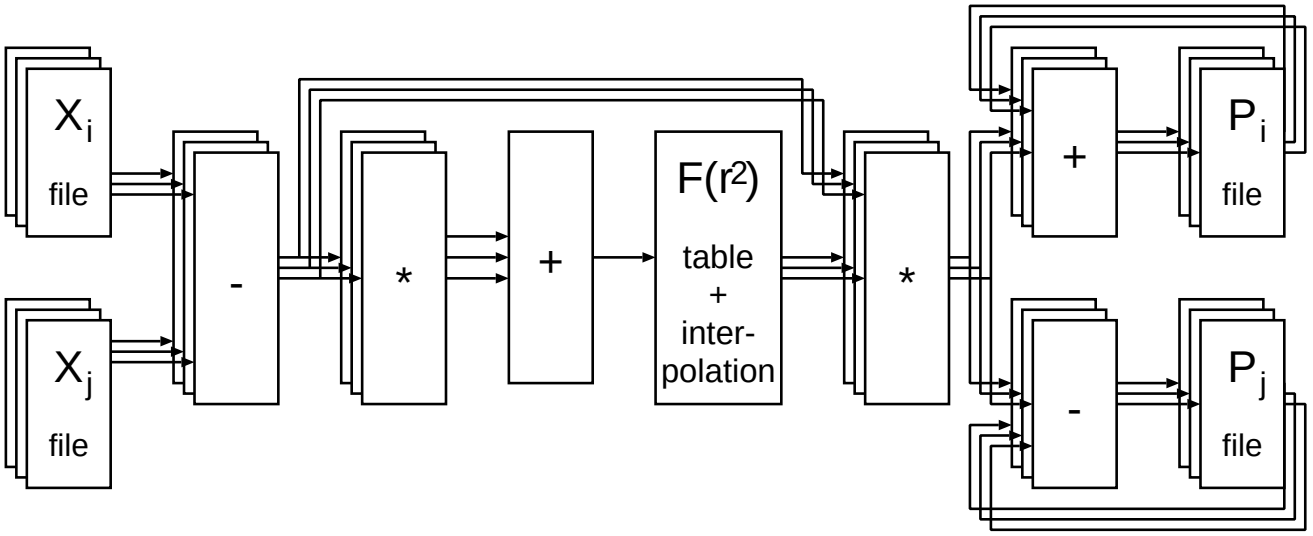

Was DMDP a great success?

— Achievement

- Faster than Cray-XMP for the cost less than 1%

— Limitations

- Not widely used, no followups
  - Rather severe limitations on the system to be simulated (practically, neutral, single-atom molecules only)
  - Analysis of the result was done in hardware: flexible analysis was impossible

Basic idea of GRAPE (GRAvity PipE)

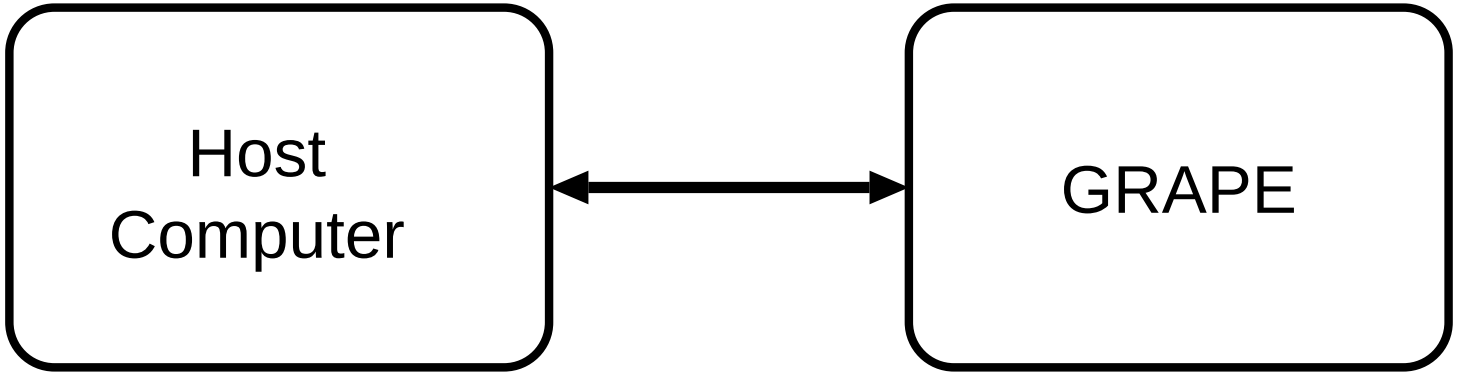

Time integration etc. Interaction calculation

Specialized hardware: interaction calculation

General-purpose computer: everything else

Specialized hardware

- Hardwired logic $\longrightarrow$ no software necessary
- specialized hardware for interaction calculation
  $\longrightarrow$ high performance

General-purpose host computer

- High-level programming languages (Fortran, C, C++...)
- Softwares developed on GPCs (even "dusty decks") can be used with little modifications
- Can be used with treecode, individual timestep, FMM ...

History of the project

GRAPE-1	(89/4 — 89/10)	120 Mflops, low accuracy
GRAPE-2	(89/8 — 90/5)	40 Mflops, high accuracy(32bit/64bit)
GRAPE-1A	(90/4 — 90/10)	240 Mflops, low accuracy
GRAPE-3	(90/9 — 91/9)	14 Gflops, high accuracy
GRAPE-2A	(91/7 — 92/5)	180 Mflops, high accuracy(32bit/64bit)
HARP-1	(92/7 — 93/3)	180 Mflops, high accuracy(32bit/64bit)
		Hermite schieme
GRAPE-3A	(92/1 — 93/7)	6 Gflops/board
		some 50 copies are used all over the world
GRAPE-4	(92/7 — 95/7)	1 Tflops, high accuracy (32bit/64bit)
		Some 10 copies of small machines
MD-GRAPE	(94/7 — 95/4)	1Gflops/chip, high accuracy(32bit/64bit)
		programmable interaction
GRAPE-5	(96/8 — 99/4)	5Gflops/chip, low accuracy

Evolution of GRAPE systems

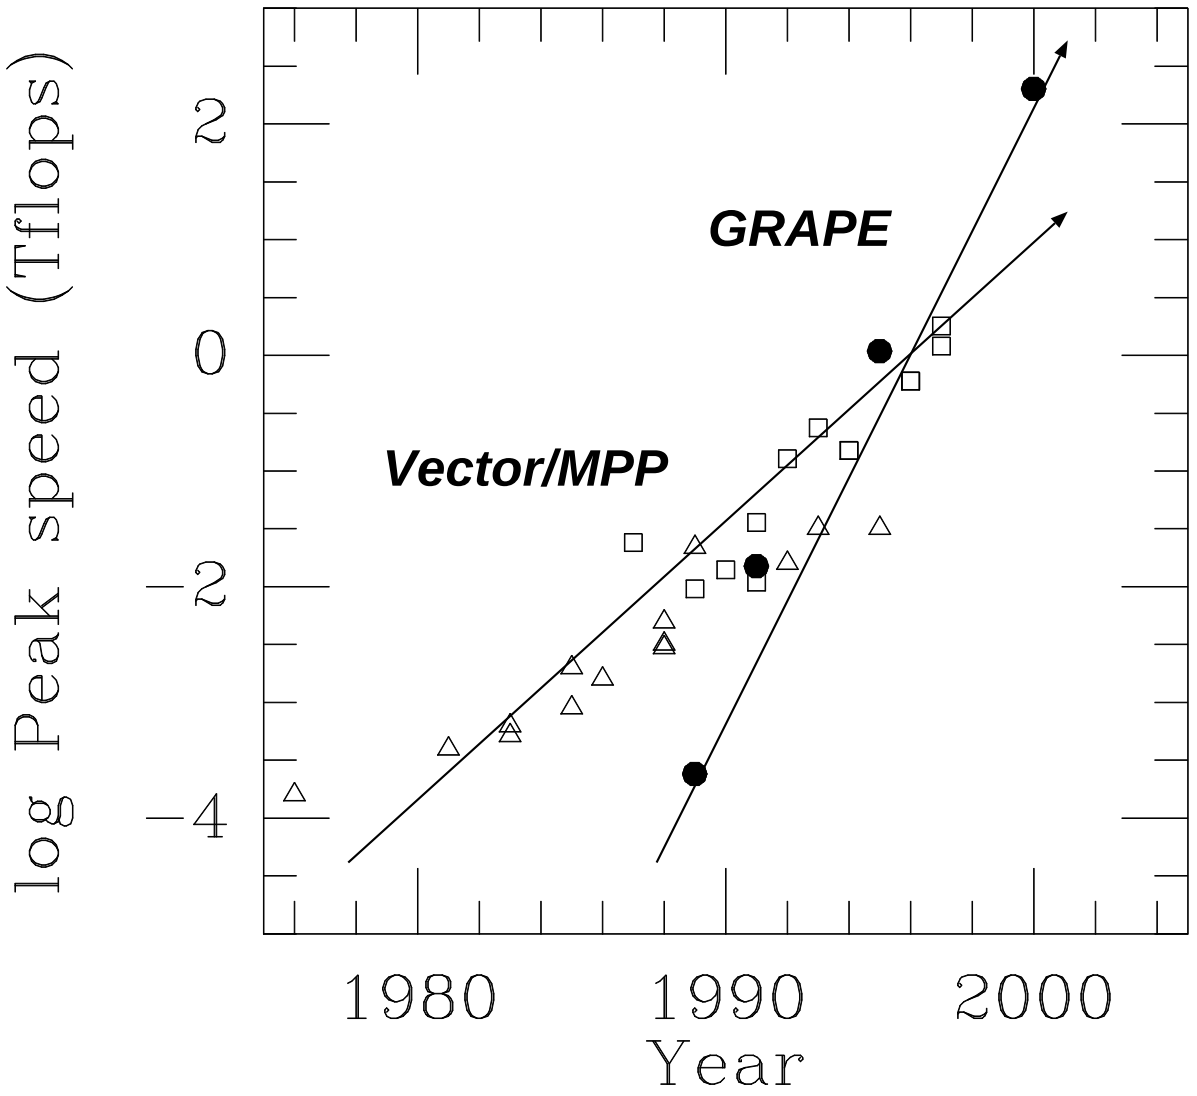

GRAPE-4

GRAPE-4 processing board

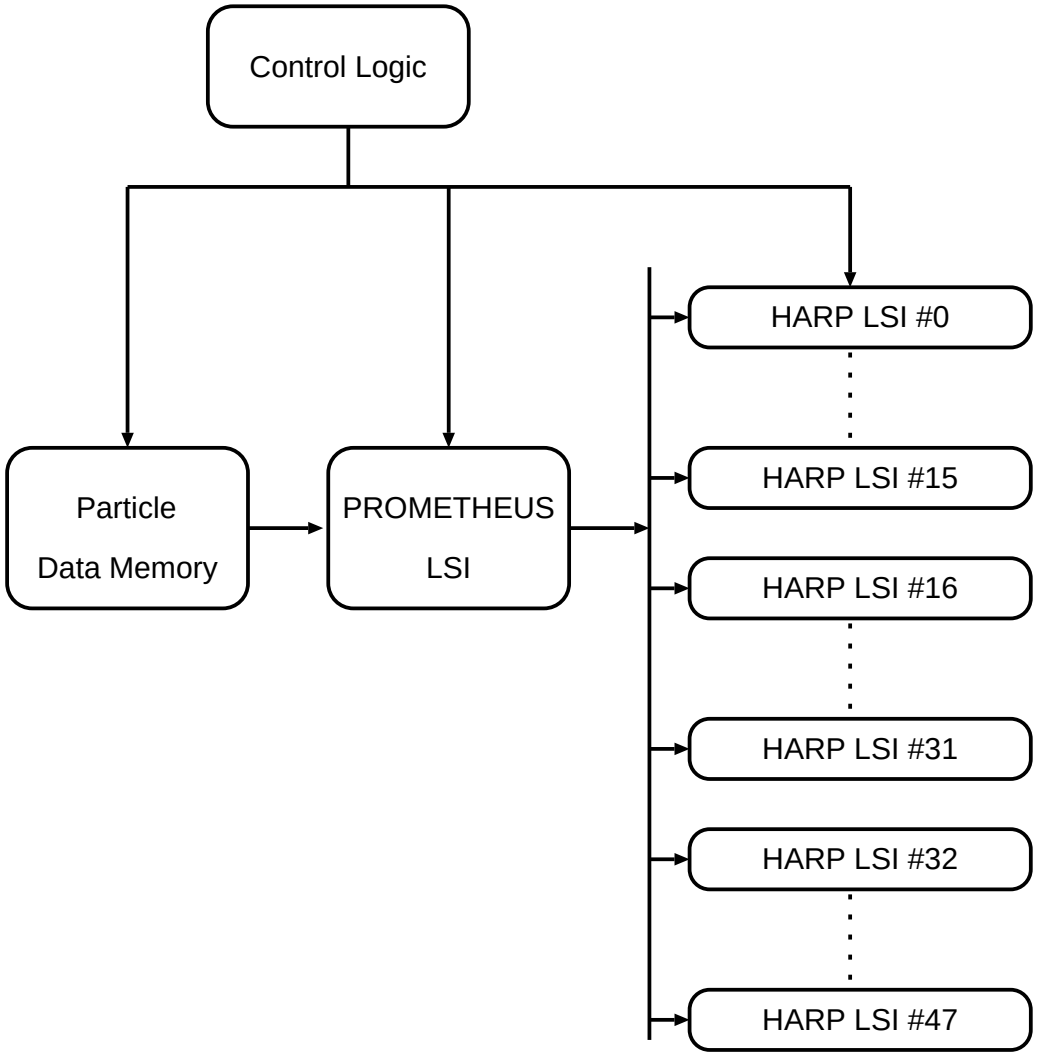

Many pipeline LSIs share one memory unit

↓

Further reduces the memory bandwidth requirement

Science with GRAPE-3/4

- Planet formation: Confirmation of the runaway growth, discovery of the slowing down and oligarchic growth.
- Globular clusters: Gravothermal oscillation, study of evolution under tidal force (Kyoto IAU, JD-15: ALL N-body results were from GRAPE-4)
- Interaction of galaxies and evolution of groups and clusters
- Structure of CDM halo
- and much more ...

GRAPE-6

— Goal

1. For gravity, achieve the speed more than 100 times faster than that of GRAPE-4 by the year 2001
2. Develop FPGA (Field Programmable Gate Array) based "programmable" machine
3. Construct a system with wide range of applications by combining the above two

Funding

Part of "Computational Science and Engineering" program under JSPS-RFTF (FY1997-2001)
Total budget: 500 M JYE

The overall structure

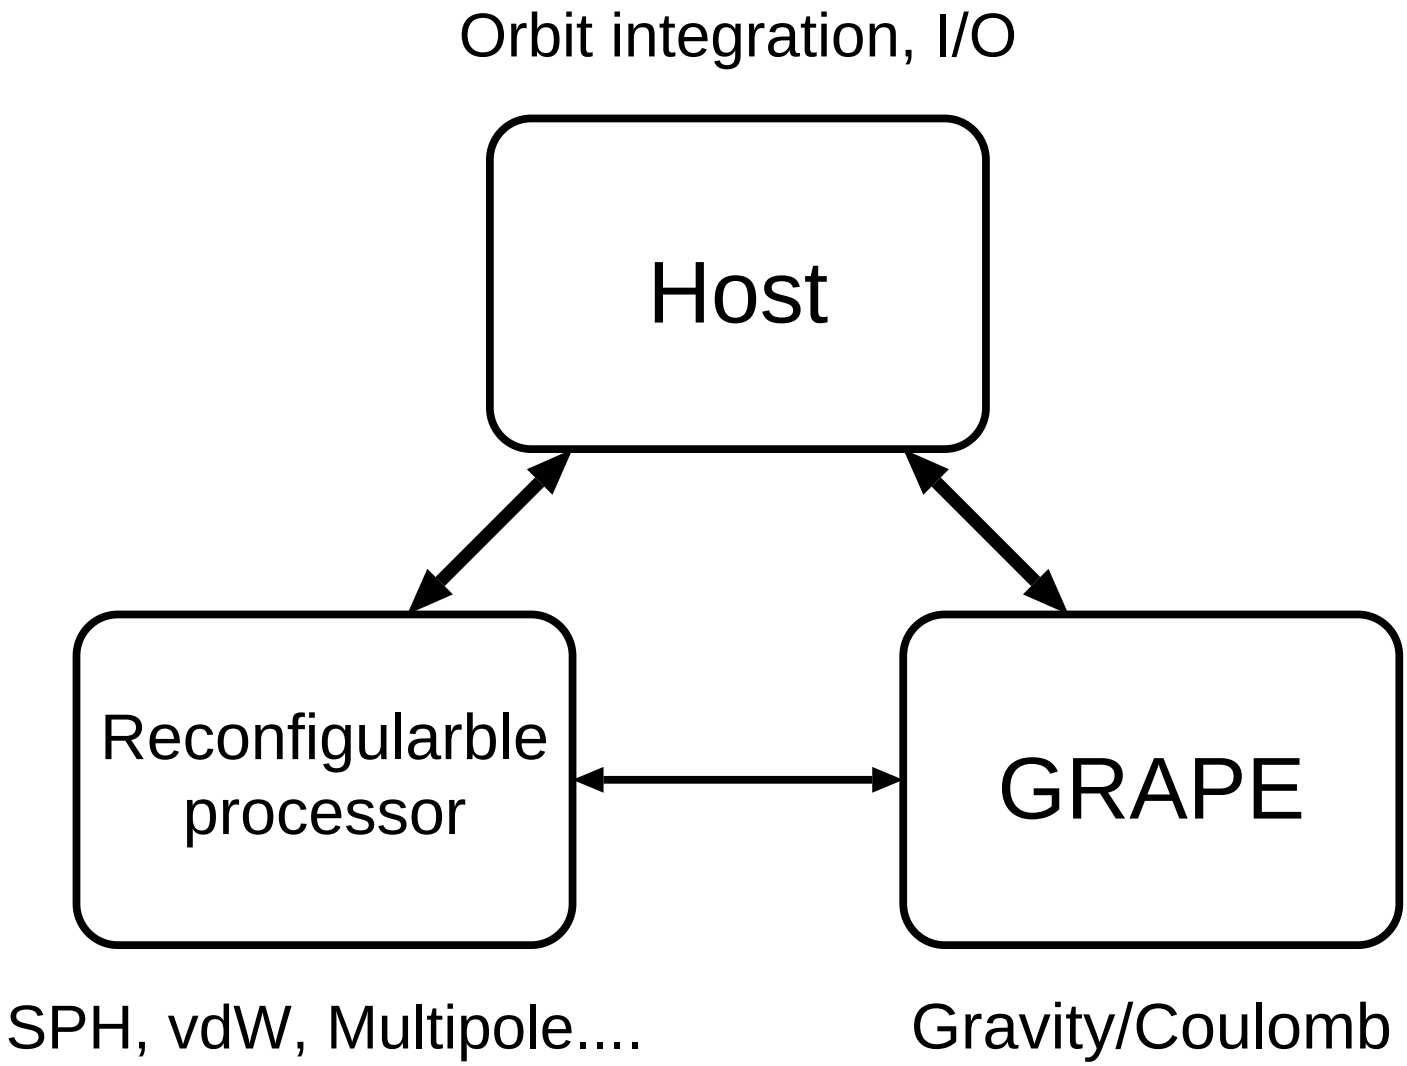

Gravity part

- Followup of GRAPE-4
- Increase the computing speed (factor of 50–70 per chip)
- Total speed exceeding 100Tflops

Overall structure

- 16 (or somewhat less...) clusters
- 16/8 processor boards per cluster
- 8 processor modules per board
- 2/4 chips per modules
- 6 pipelines per chip
- 60 arithmetic units per pipeline

In total,1.5×10^6 Arithmetic units operating at 90 MHz = 135 Tflops

GRAPE-6 overall architecture

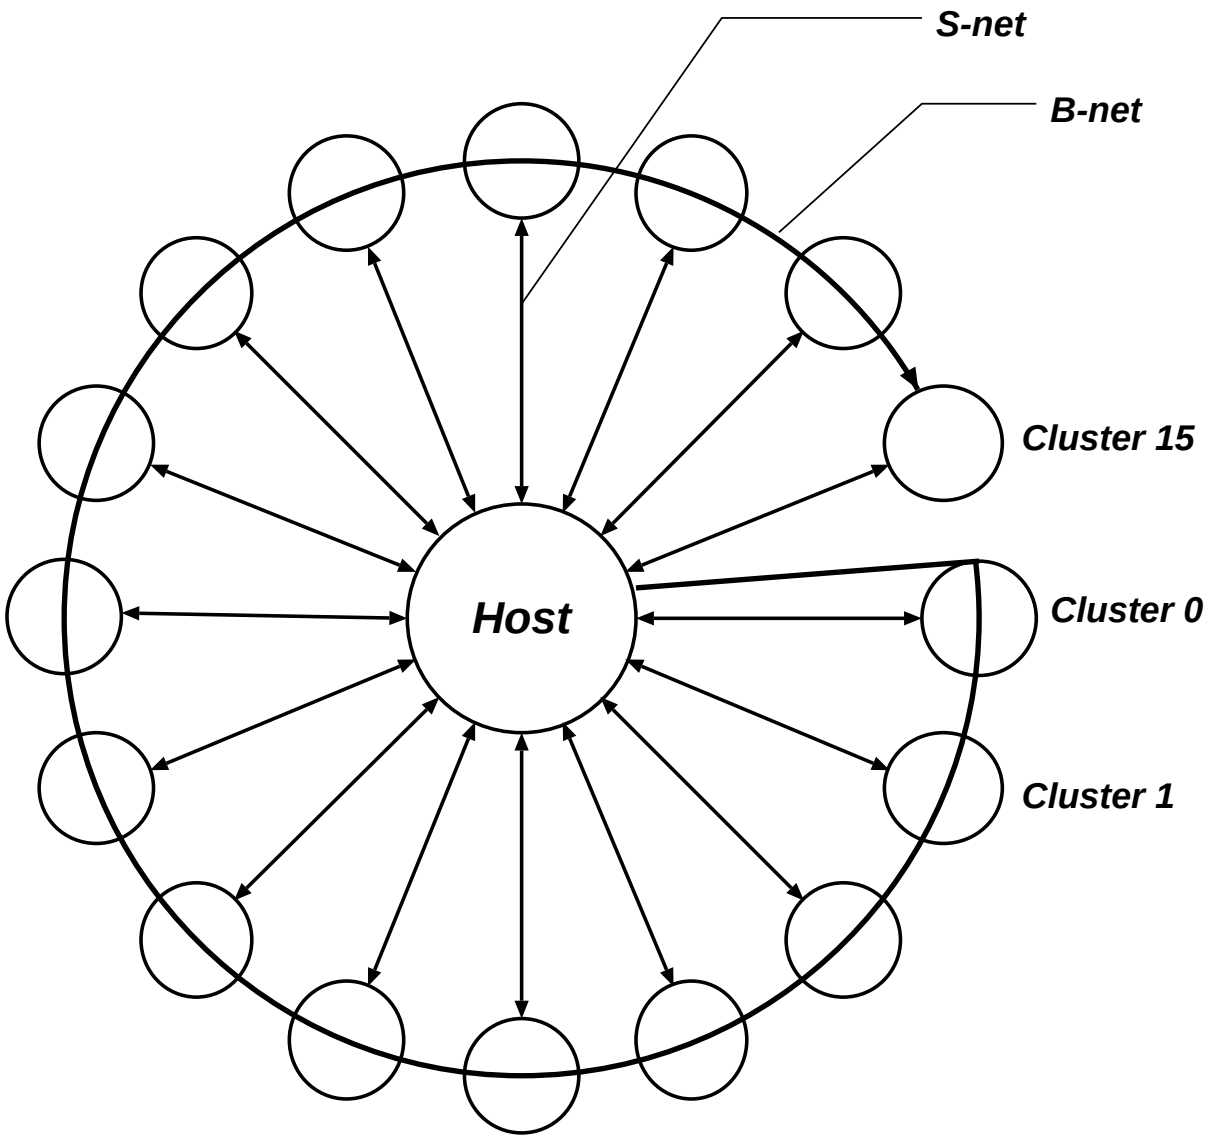

Difference from GRAPE-4

- Higher peak speed (100×)
- Higher communication speed (target: 50×)
- "System on a chip" approach
- (Hopefully) more reliable — ECC/parity check

Processor LSI

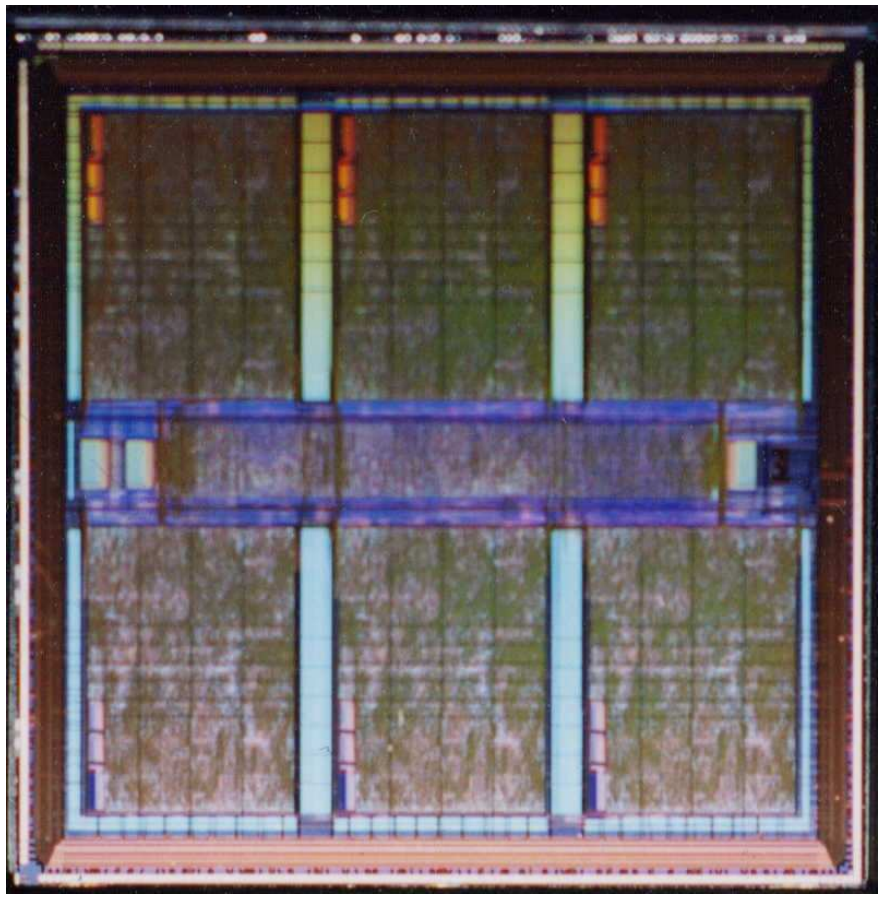

- 0.25 μm design (Toshiba TC-240, 1.8M gates)
- 6 pipelines in a chip
- 30–35 Gflops (equivalent) per chip

GRAPE-6 processor module

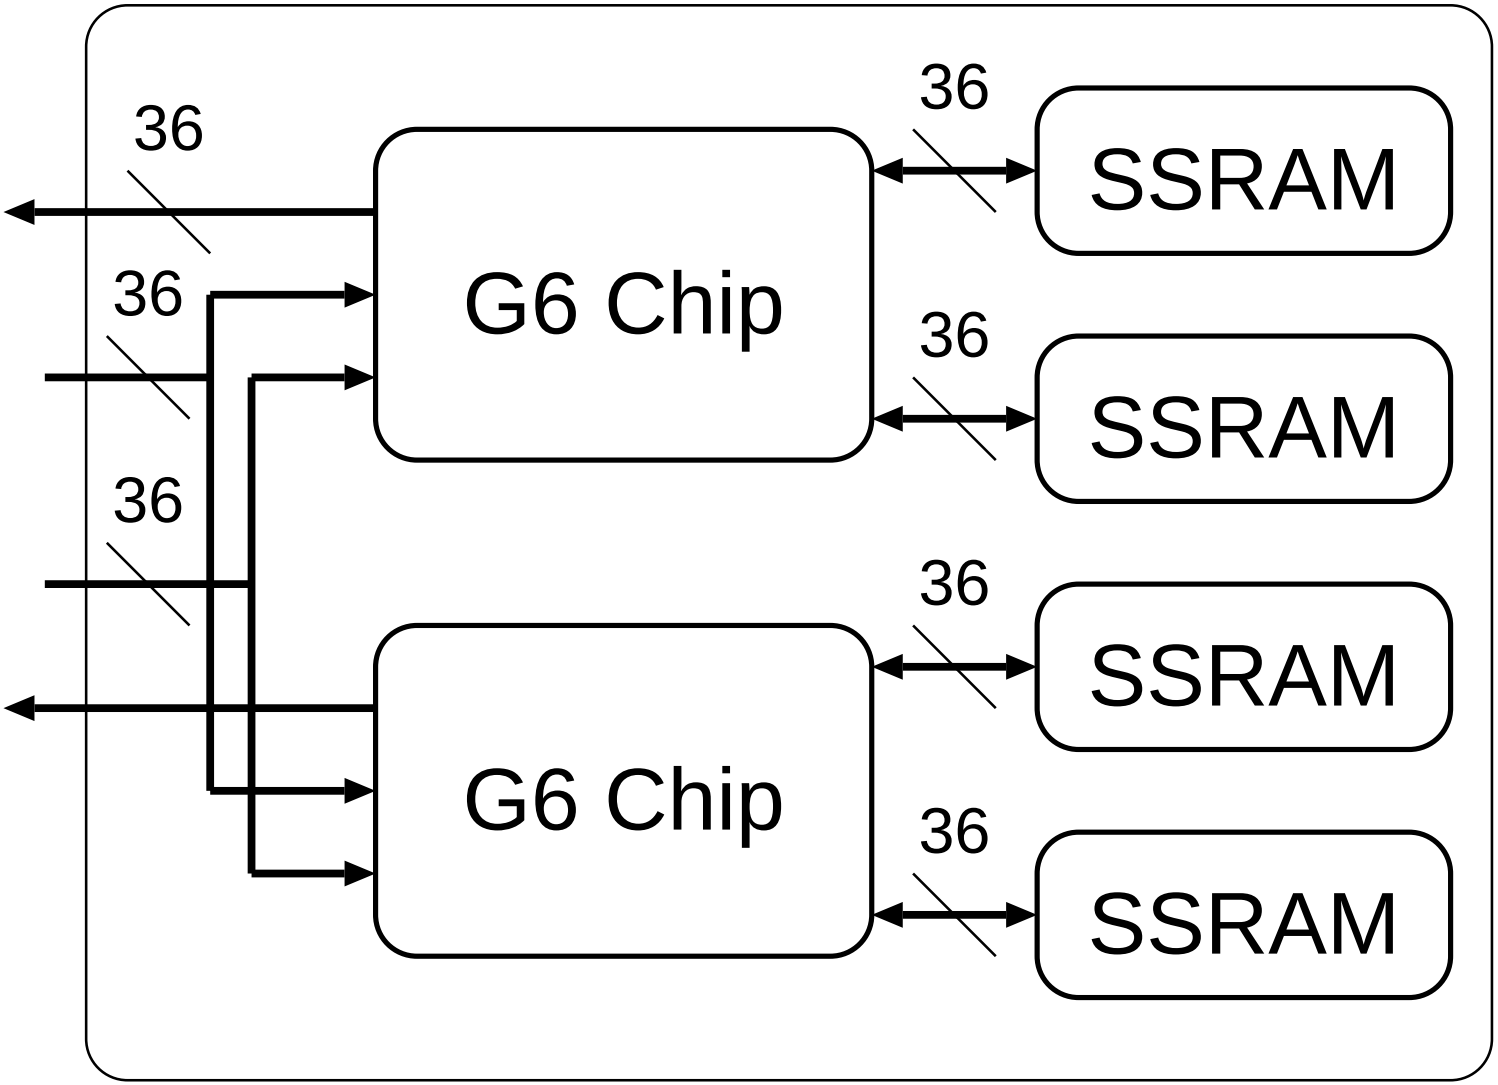

GRAPE-6 processor module

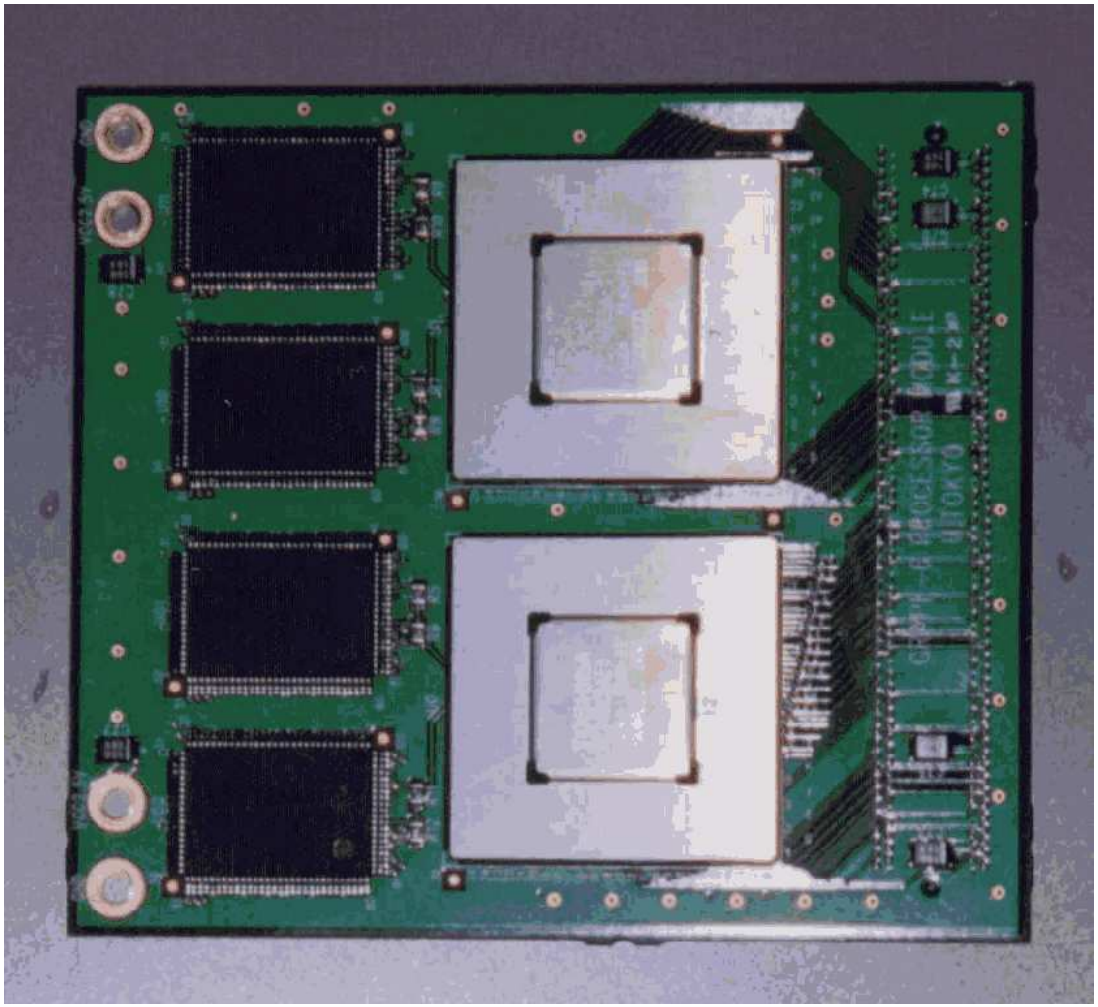

- Two processor chips + Memory chips.
- Integrates all functions of GRAPE-4 processor board.

The GRAPE-6 Processor board

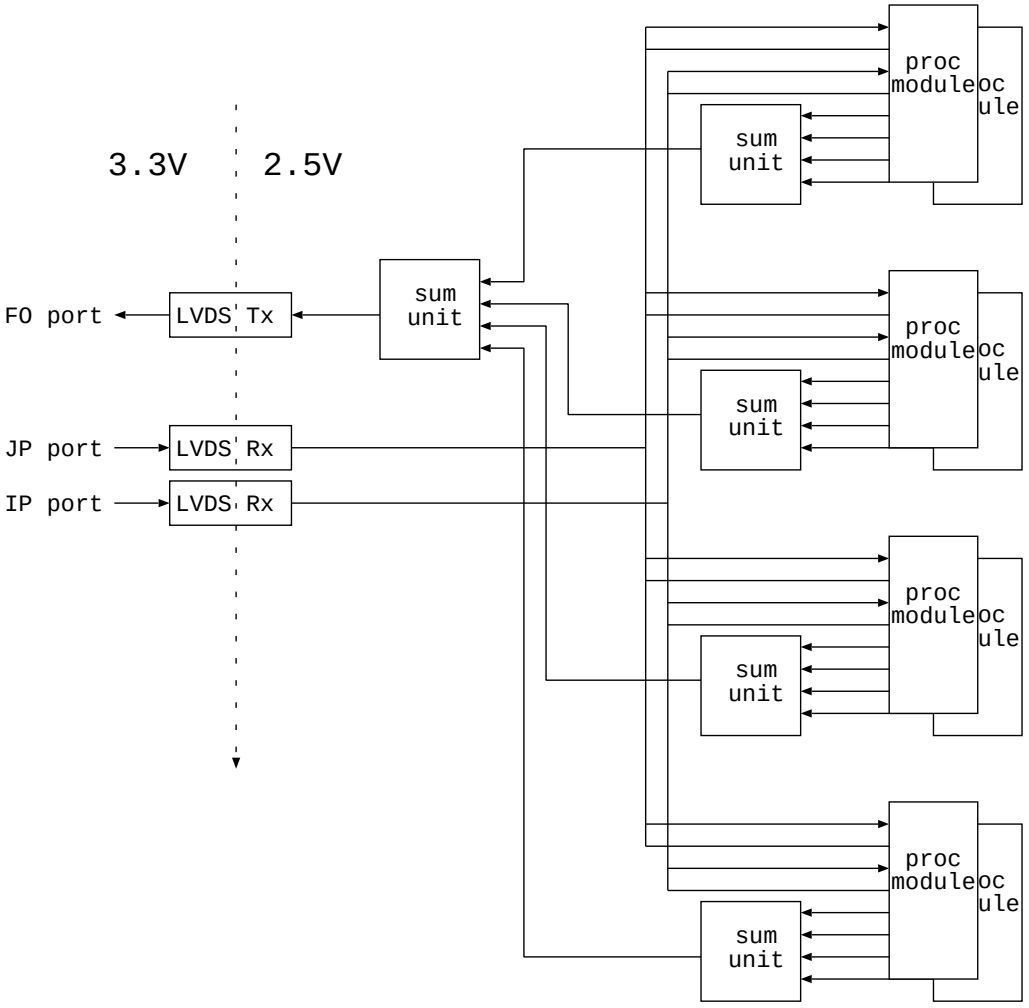

- Single board with 16-chips
- Semi-serial (LVDS) interface (350MHz clock, 4 wires)
- Two input ports and one output port
- Result from 16 chips are summed by a reduction tree

The GRAPE-6 processor board

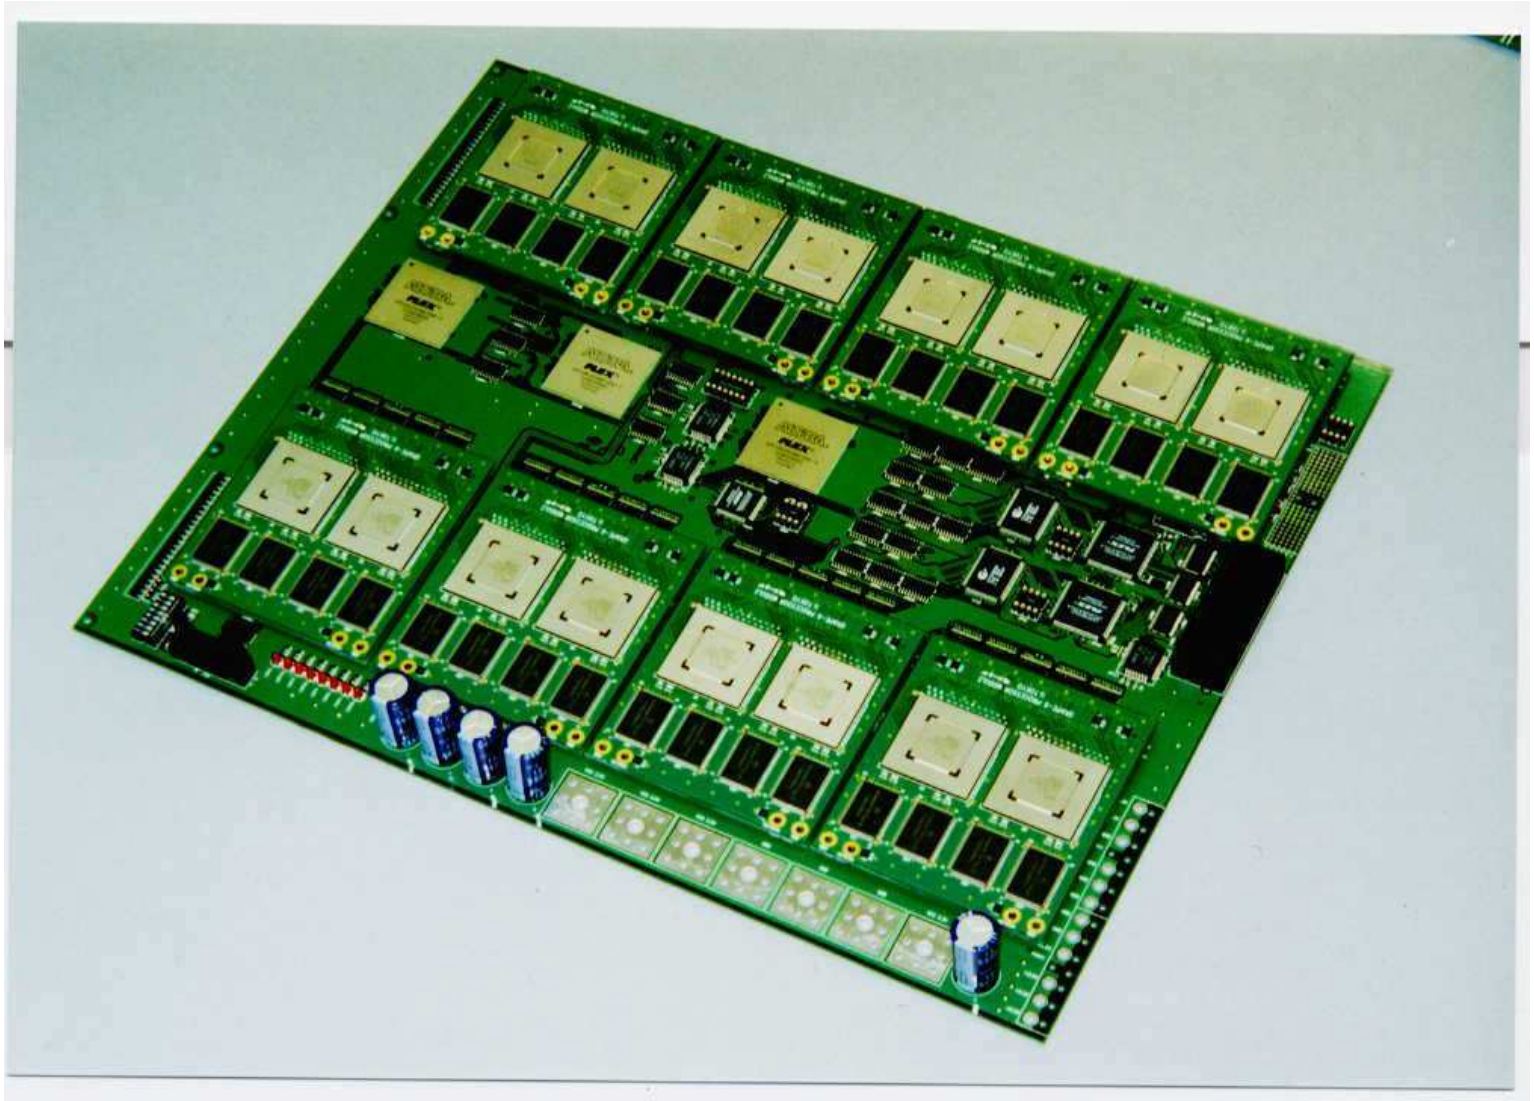

GRAPE-6 cabinet and 4-board system

Present status

- 4-board, 2-Teraflops system is now up and running.
- Current clock: 88 MHz
- Number of chips: 64

FPGA-based machine

GRAPE accelerates long-range interaction

↓

short-range forces become the next bottleneck

Example of the bottleneck:

SPH interaction, vdW forces in MD calculation etc

A multi-purpose machine for these types of interaction =
PROGRAPE (PROgrammable GRAPE)

Structure of PROGRAPE

Same basic idea as "traditional" GRAPE

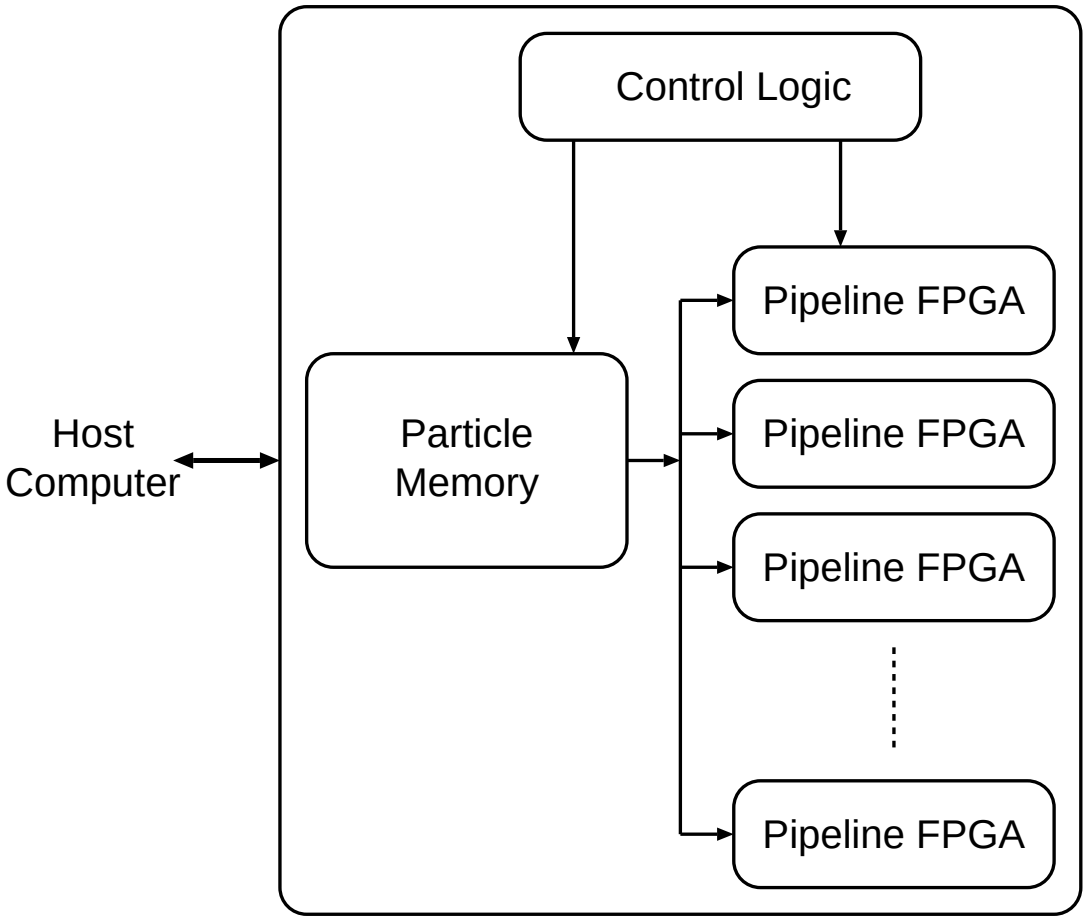

PROGRAPE-1

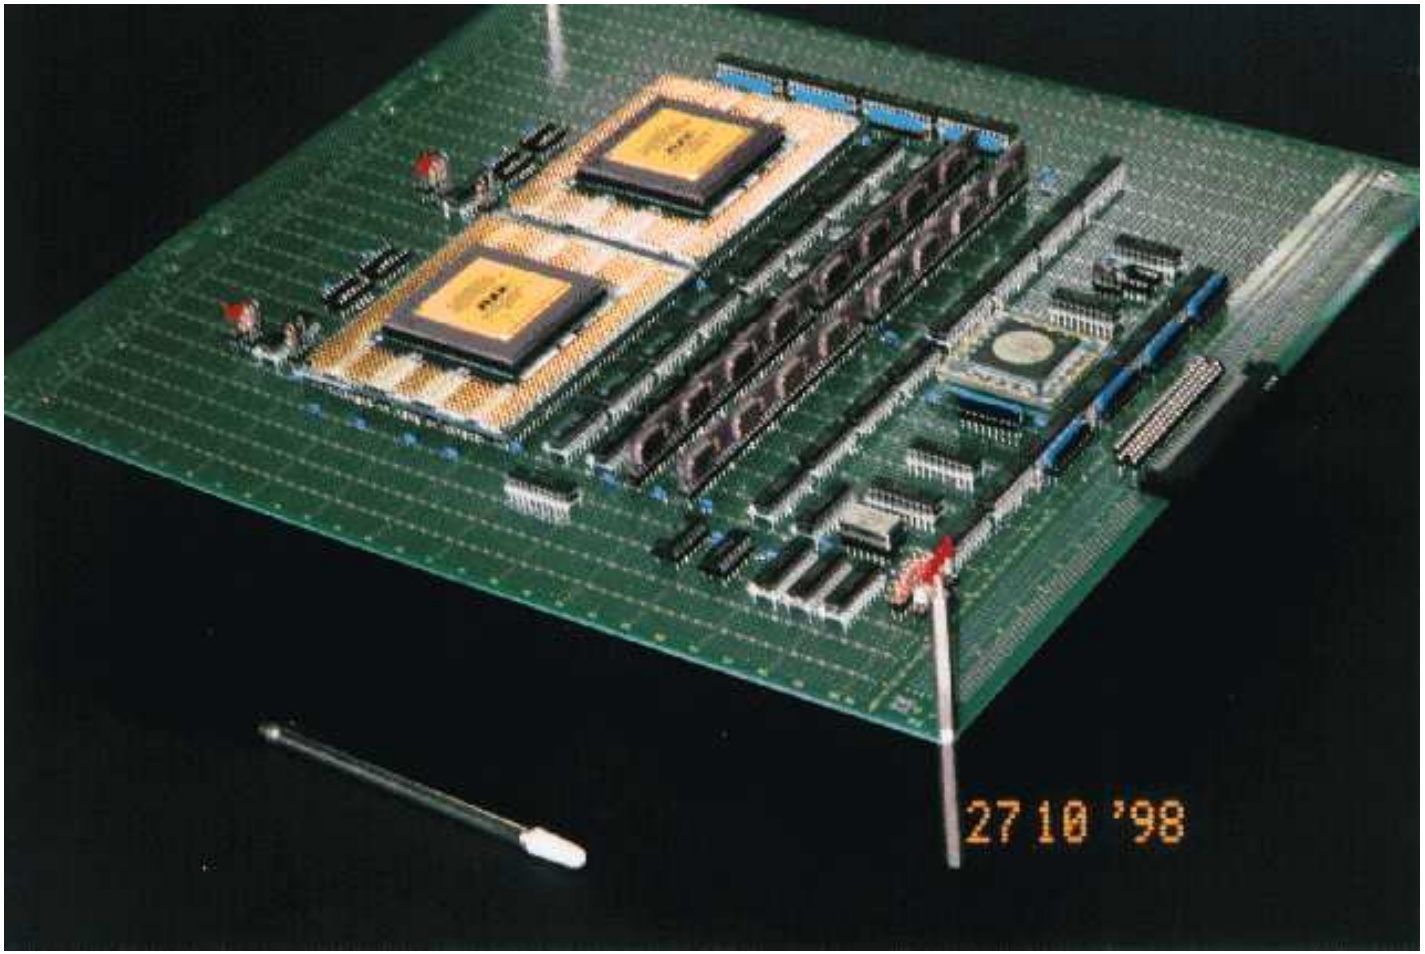

MDM project

- Project to build a 150-Tflops machine for classical MD calculation
- Ebisuzaki, Narumi, Susukida and many others in RIKEN
- Project started in 1995
- Hardware almost completed

Summary

- Special-purpose computers such as GRAPE have very large advantage over GPCs.
- The advantage has been, and will be, increasing *exponentially* in time.
- GRAPE-4 did make use of such potential advantage, and has produced many new scientific results, and will produce more.
- GRAPE-6 will be available in year 2001, and will be *100 times faster* than GRAPE-4.