GRAPE-DR
— Next-Generation GRAPE Project

Jun Makino

University of Tokyo

Talk overview

- **GRAPE Project**
- **Science with GRAPEs**
- **Next Generation GRAPE — the GRAPE-DR**

GRAPE project

GOAL:

Design and build specialized hardware for simulation of stellar systems.

Rational:

***N*-body simulation is compute-intensive.**

With specialized hardwares, you can do larger simulations (better resolution) for same amount of money.

GRAPE-6	**(2002, 64 TF)**	**4M$**
ASCI White	**(2001, 12 TF)**	**200M$**
ASCI Q	**(2002, 30 TF)**	**200M$**
Earth Simulator	**(2002, 40 TF)**	**300M$**

(running cost is roughly proportional to the price)

Basic idea of GRAPE

Special-purpose hardware for force calculation
General-purpose host for all other calculation

$$\frac{d^2\mathbf{x}_i}{dt^2} = \sum_{j \neq i} -G m_j \frac{\mathbf{x}_i - \mathbf{x}_j}{|\mathbf{x}_i - \mathbf{x}_j|^3}$$

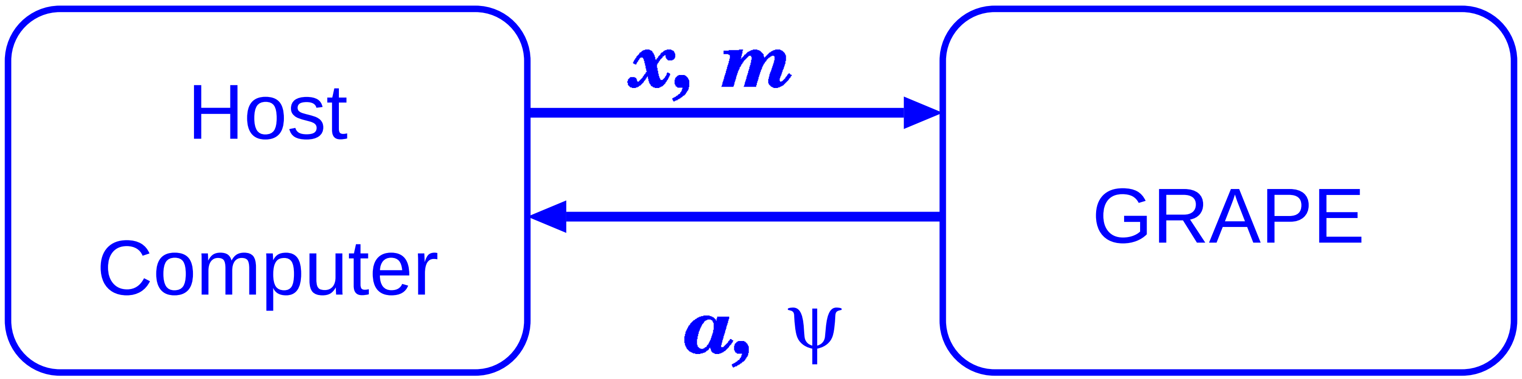

Time integration, IO, etc

Force calculation

Flexibility

High performance

GRAPE pipeline processor

GRAPE = GRAvity PipE

Original proposal: Chikada (1988)

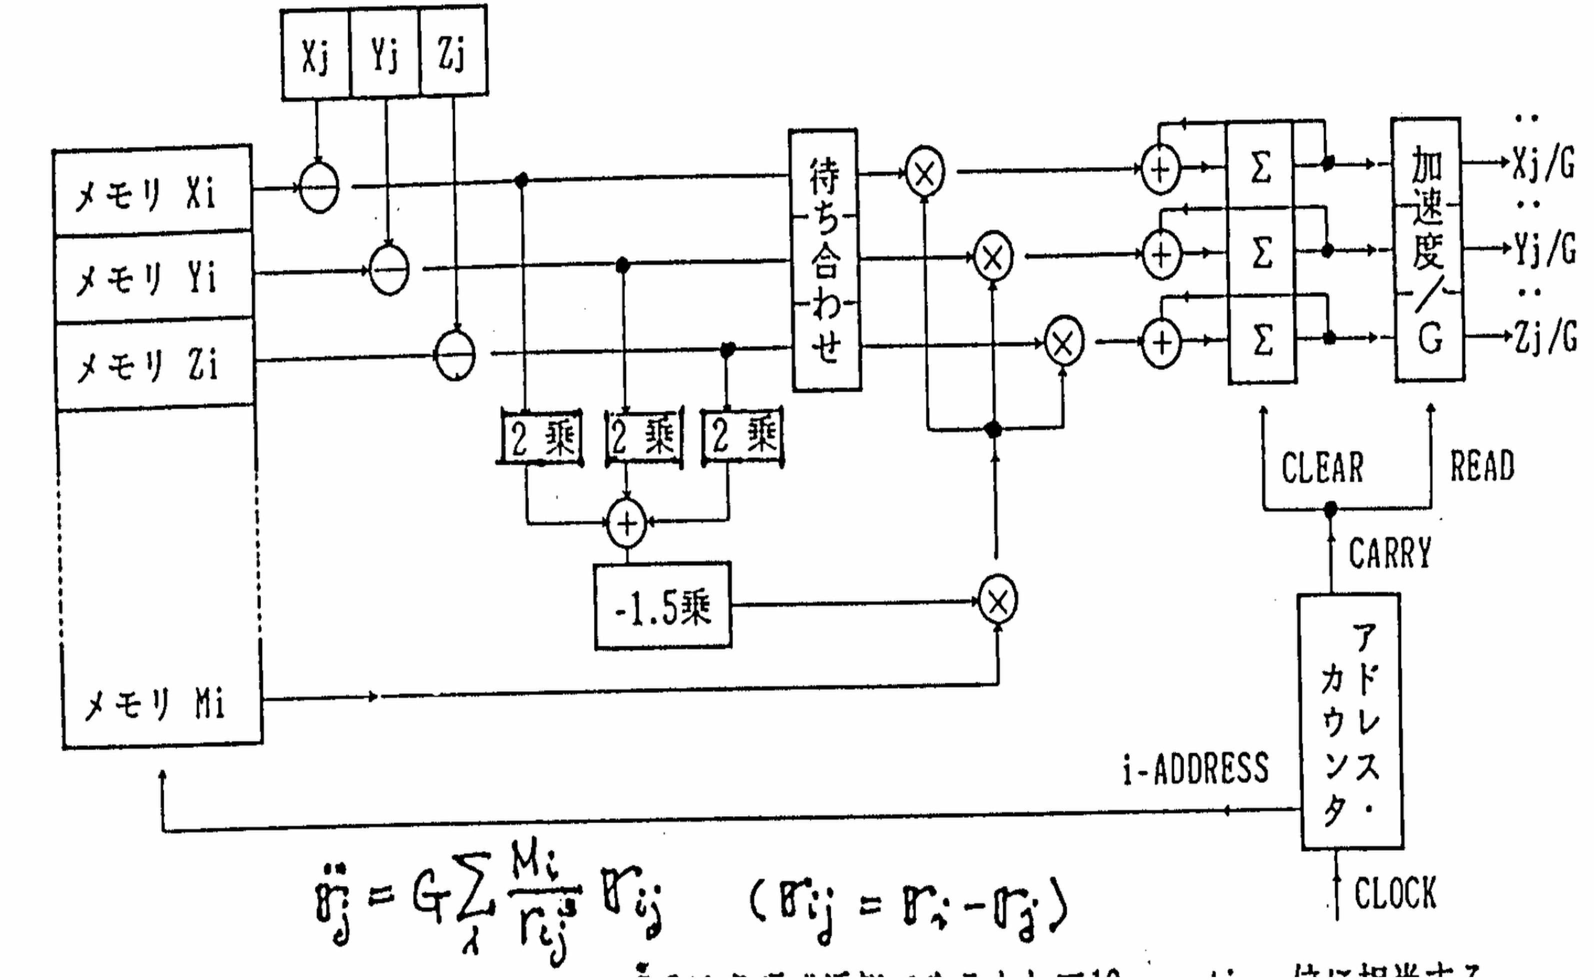

GRAPE pipeline processor(2)

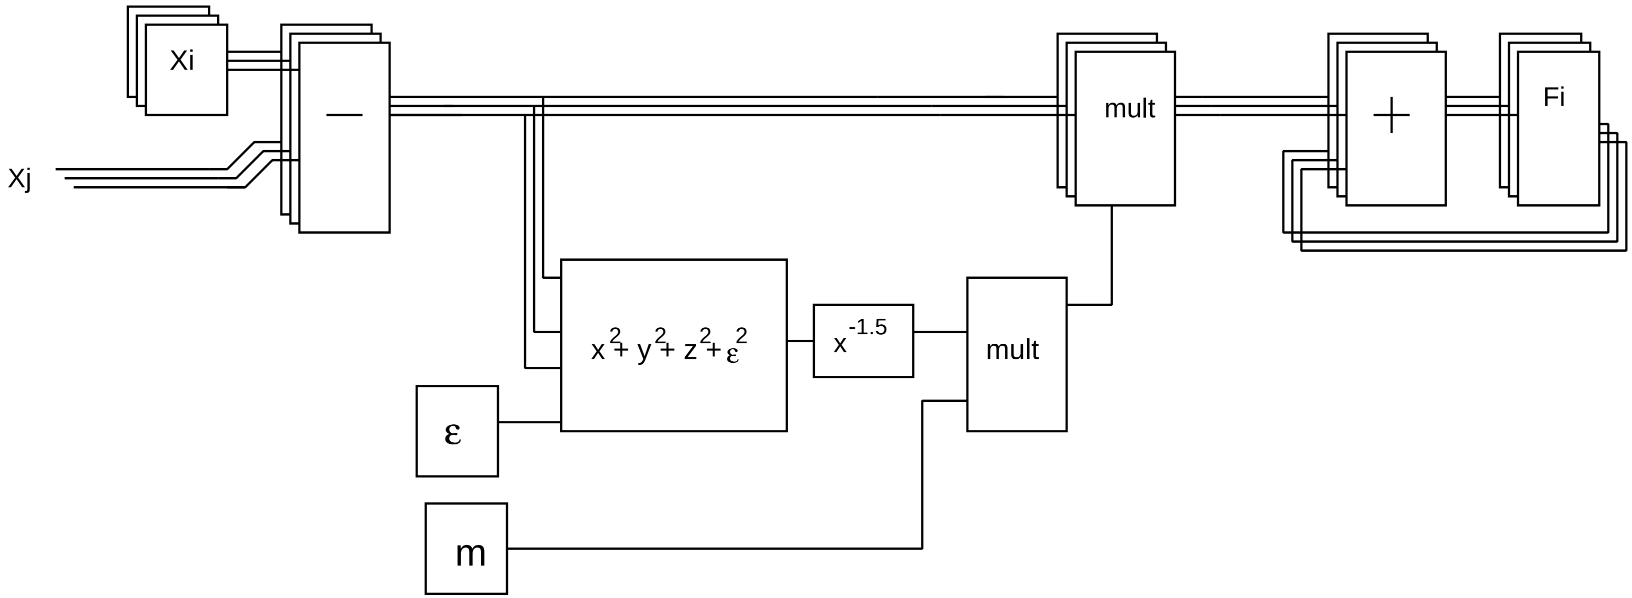

$$\mathbf{F}_i = \sum_{j \neq i} m_j \frac{\mathbf{x}_j - \mathbf{x}_i}{(|\mathbf{x}_j - \mathbf{x}_i|^2 + \epsilon^2)^{3/2}}$$

GRAPE hardwares

- **Low-accuracy machines**
  - **for collisionless systems**
  - **Very high performance for low cost**
- **High-accuracy machines**
  - **for collisional systems**
- **Machines for non-1/r force laws**
  - **Molecular Dynamics**
  - **SPH**

GRAPE History(1)

Low accuracy machines

Machine	Year	Peak Speed	Notes
GRAPE-1	1989	120 Mflops	First GRAPE
GRAPE-3	1991	14 Gflops	Custom LSI
GRAPE-5	1998	40 Gflops	2 pipes/chip

High accuracy machines

Machine	Year	Peak Speed	Notes
GRAPE-2	1990	40 Mflops	High accuracy
HARP-1	1993	180 Mflops	Hermite scheme
GRAPE-4	1995	1.1 Tflops	1728 chips
GRAPE-6	2002	64 Tflops	2048 chips

GRAPE History(2)

Non $1/r$ force laws

Machine	Year	Peak Speed	Notes
GRAPE-2A	1992	180 Mflops	High accuracy
MD-GRAPE	1995	4.2 Gflops	Custom LSI
MDM	2002	75 Tflops	RIKEN project
PE	2005?	1Pflops	RIKEN project

Evolution of GRAPE systems

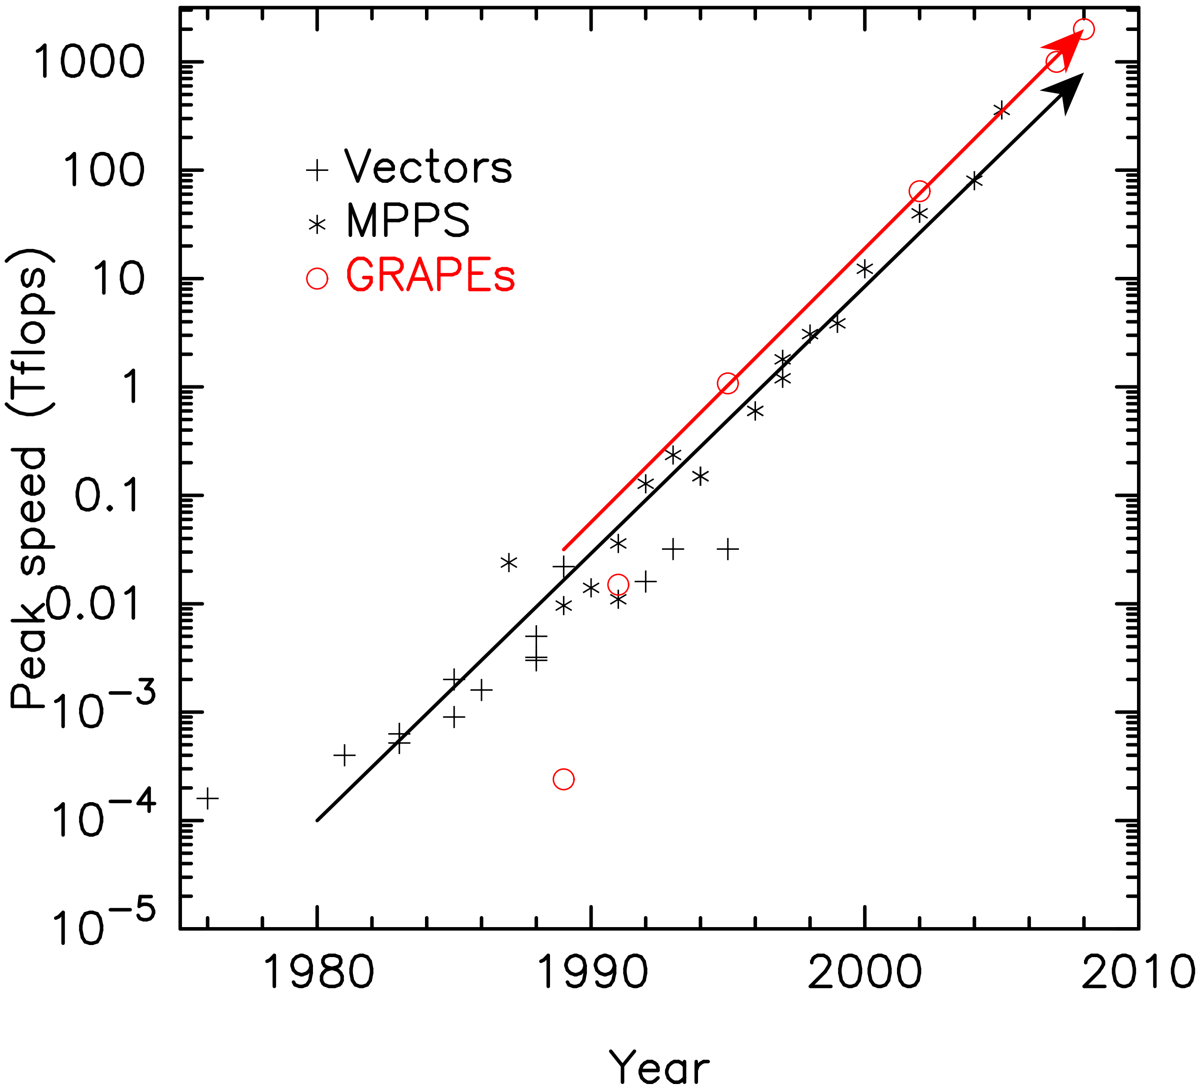

Faster than the fastest general-purpose computers for a decade.

The cost of general-purpose machines has increased dramatically in 1990s ($30M → $400M)

The grant for GRAPE did not...

GRAPE-6

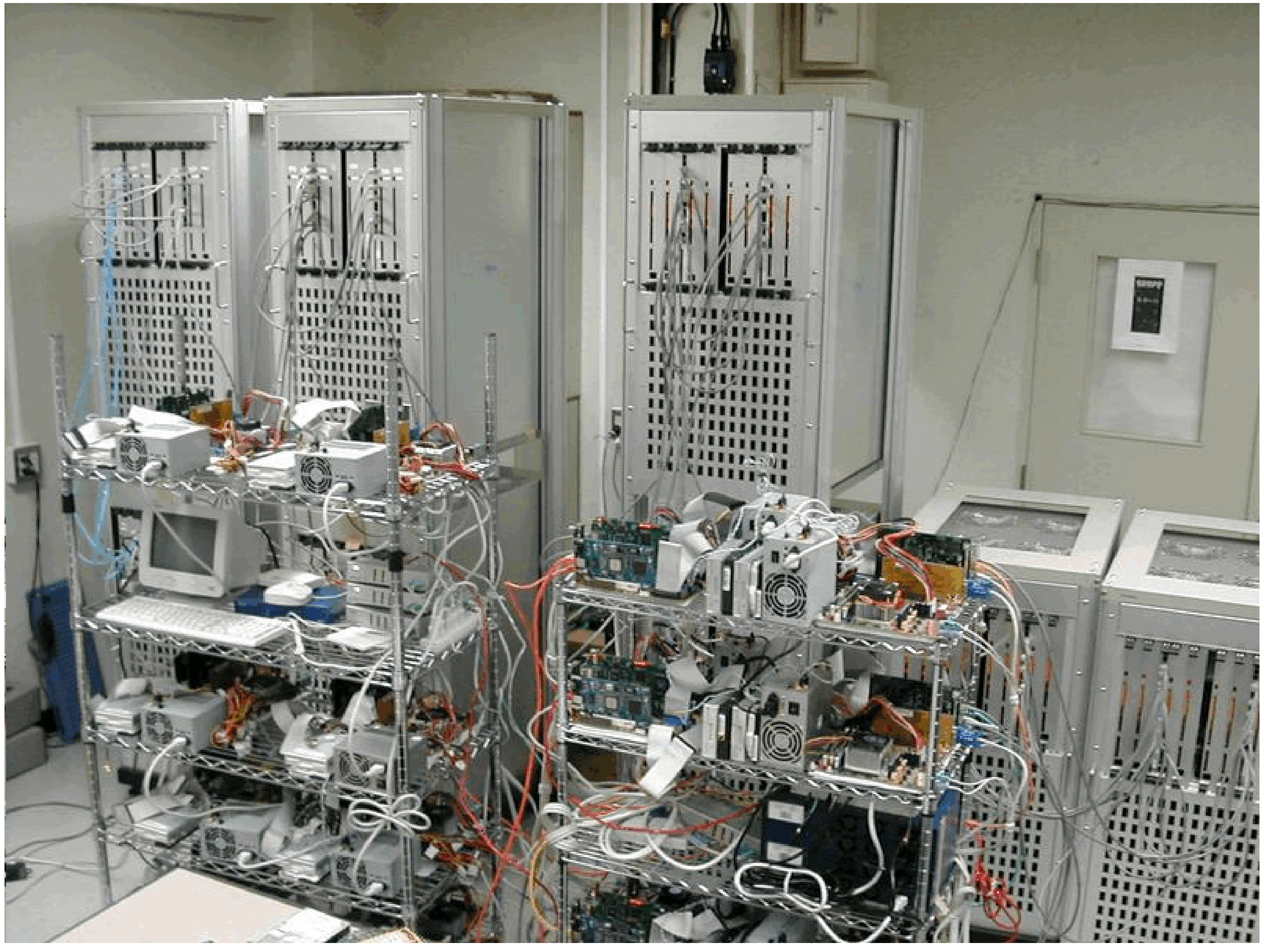

64 Tflops peak speed

Measured best performance: 35.3 Tflops

Three Gordon-Bell Prizes (2000, 2001, 2003)

The Earth Simulator

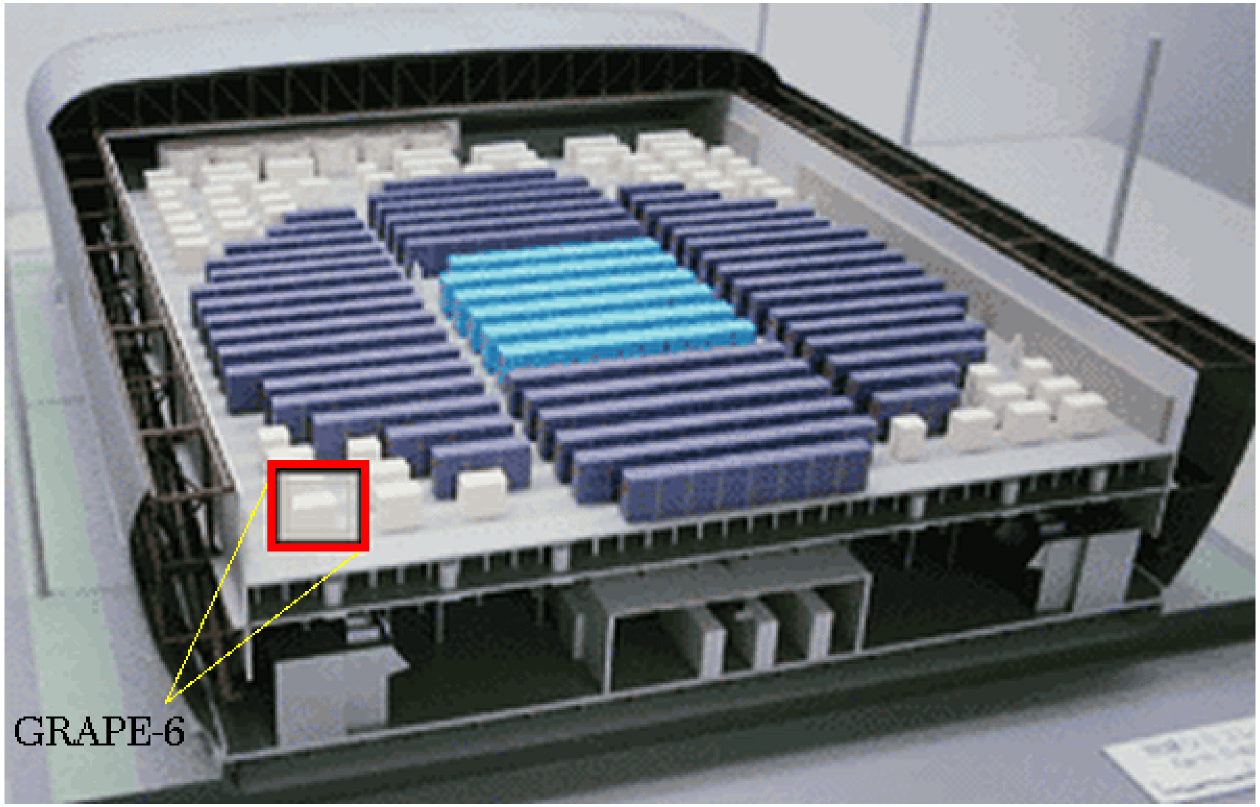

40 Tflops
peak speed

Measured
best
performance:
26 Tflops

Three
Gordon-Bell
Prizes
(2002-2004)

GRAPE-6 Processor pipeline

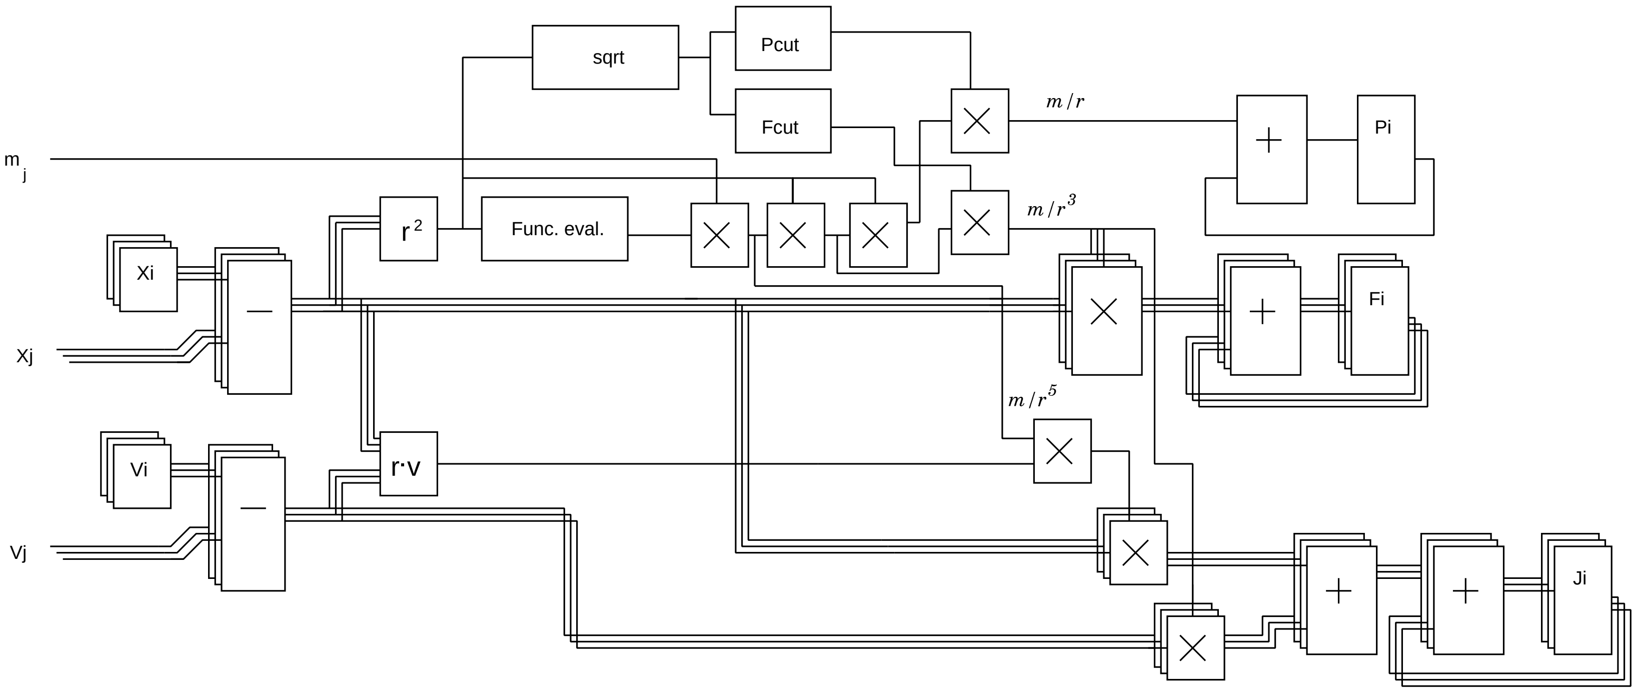

Calculates gravitational force, its first time derivative and potential.

GRAPE-6 Processor Chip

- 0.25 μm custom design (Toshiba TC240E, 1.8M gates)
- 2.5 V power supply, 90 MHz operation
- 6 pipelines in a chip
- 31 Gflops (equivalent) per chip
- system clock: 22.5 MHz

GRAPE-6 Processor Chip

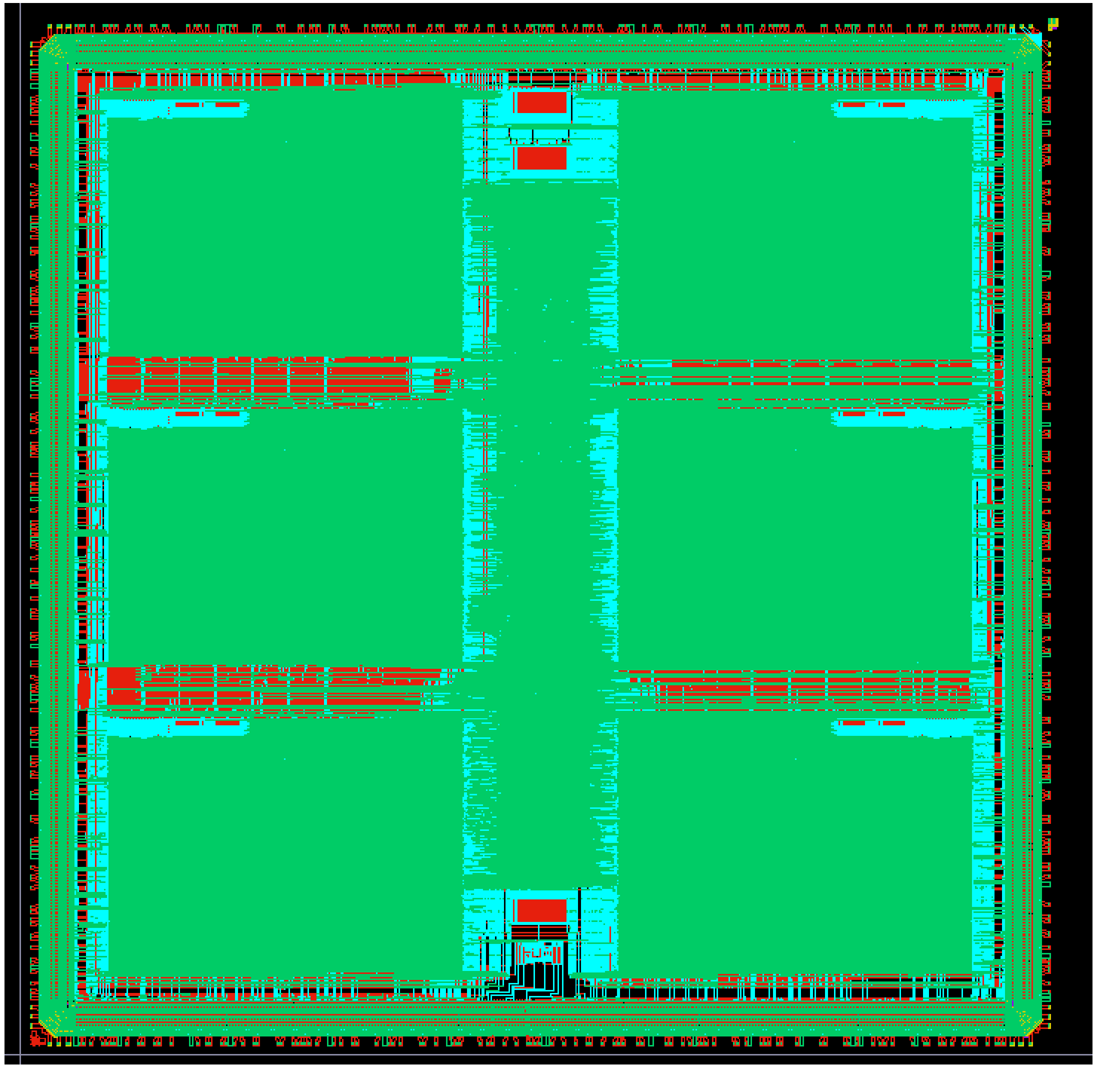

GRAPE-6 vs Intel P4

	G6	P4
Design Rule (μm)	0.25	0.09 (Pr)
No. transistors	7M	125M
Clock freq.	90M	3.8G
Power Consumption	15W	$\sim$ 100W
No. ALUs	$\sim$ 400	2 (SSE2)
Peak speed(GF)	31	7.6

- 3,500 times more efficient use of transistors
- 30 times less heat per operation
- 6 years older technology

GRAPE-6 processor module

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

The GRAPE-6 Processor board

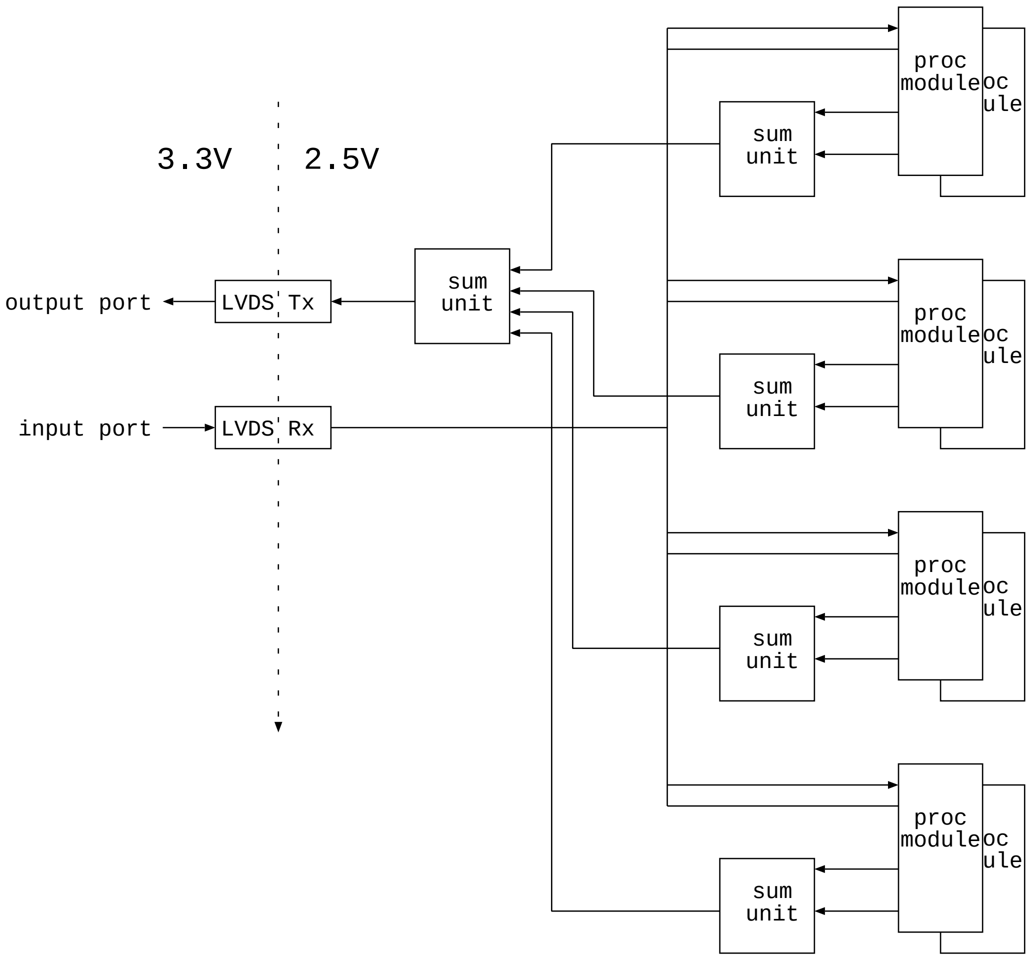

- Single board with 32 chips
- Semi-serial (LVDS) interface (350MHz clock, 4 wires)
- Two input ports and one output port
- Result from 32 chips are summed by a reduction tree

The GRAPE-6 processor board

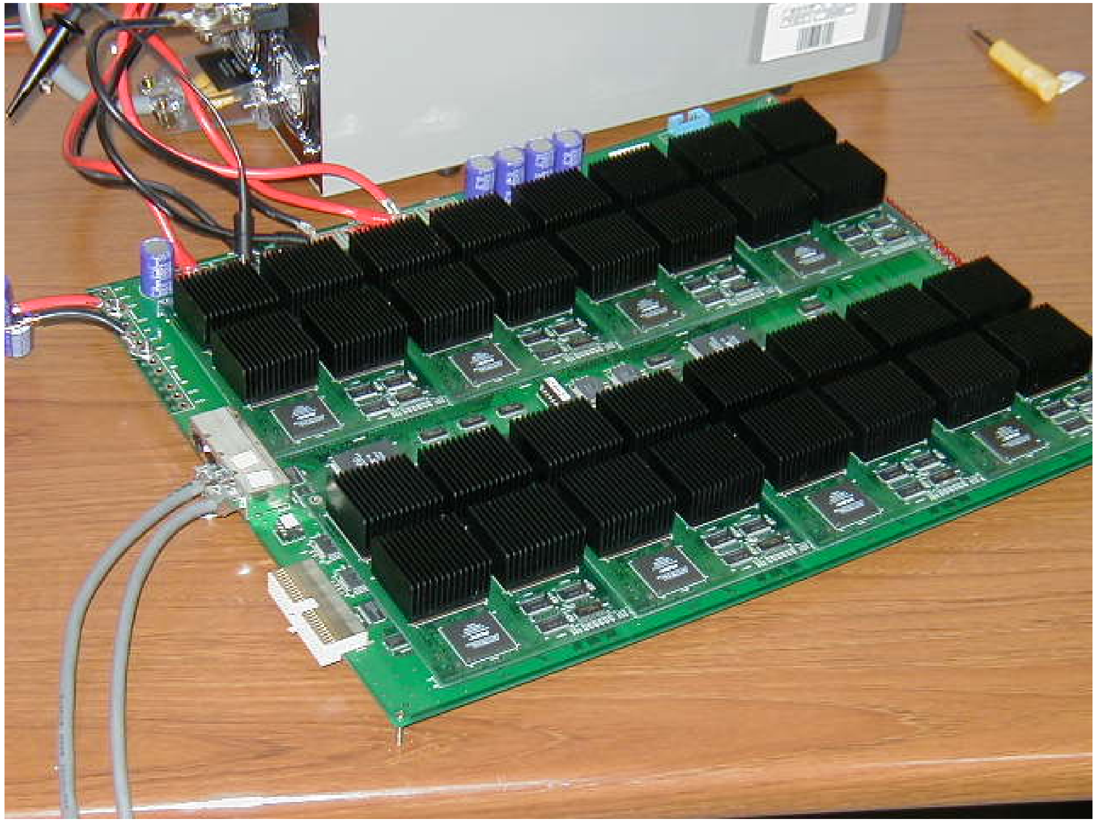

GRAPE-6

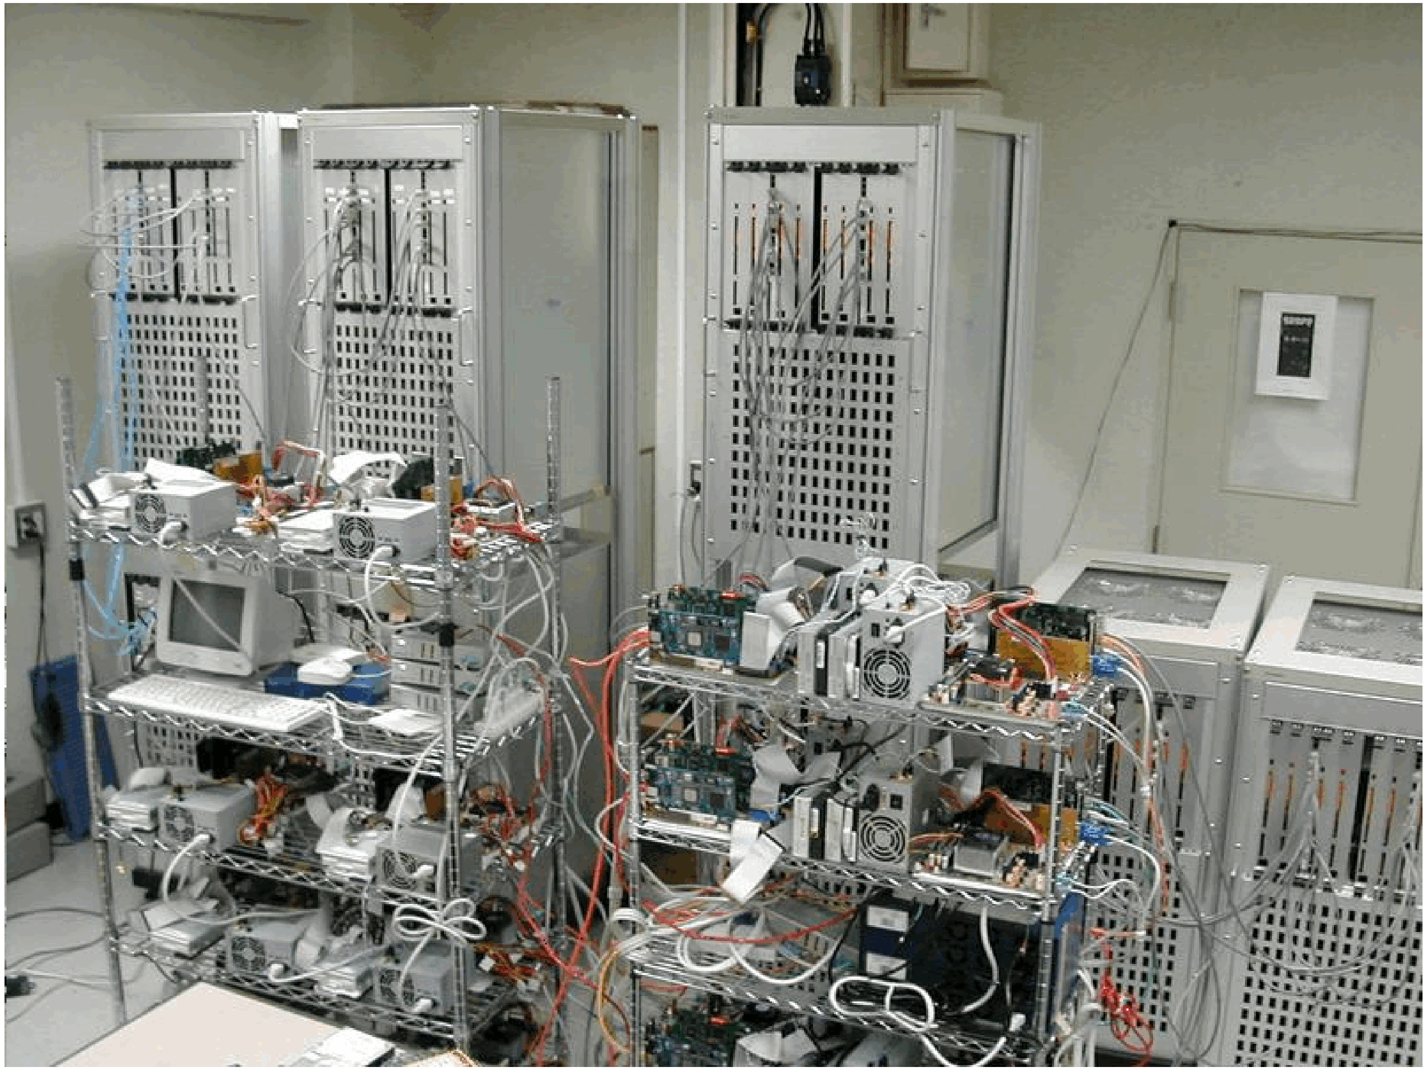

GRAPE-6A

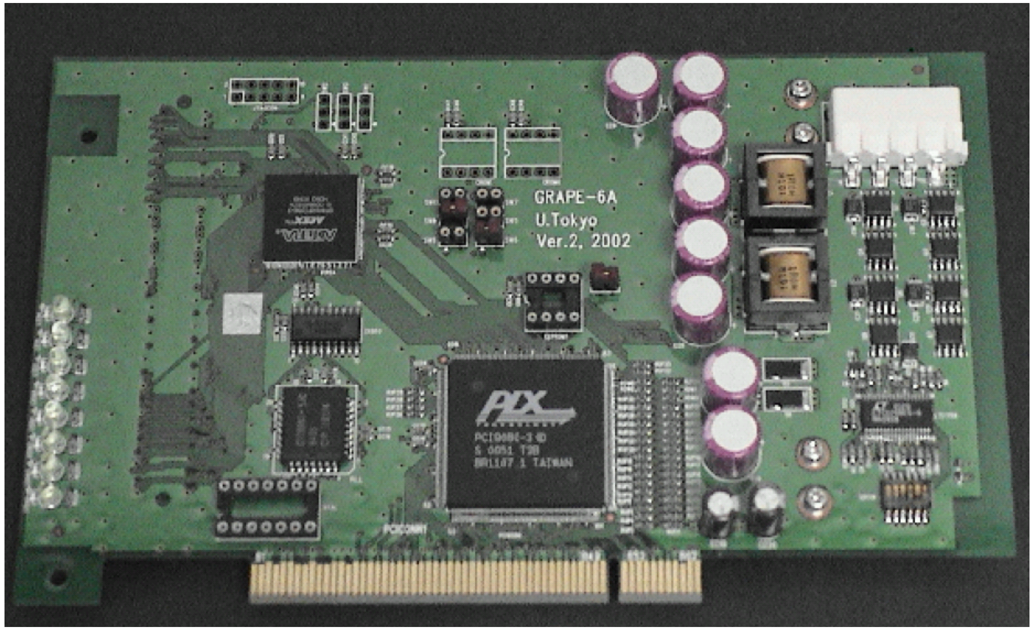

- Single PCI card with 4 chips (130 Gflops peak)
- Commercial version from Hamamatsu Metrix (www.metrix.co.jp)

- Ideal for parallel treecode/P^3M
- $6K within Japan

GRAPE-6s outside Tokyo University

Not a complete list.... Sorry if your institute is not listed.

- **NAOJ (10? GRAPE-6 and 10? GRAPE-6A)**
- **TIT (2? G6A)**
- **Gunma Observatory (1 G6)**
- **NIFS (1 G6)**
- **Tsukuba [8 G6 + 256 G6A(in 2005-2006)]**
- **Rochester (32 G6A)**
- **AMNH (3 G6)**
- **Drexel (1 G6)**
- **Heidelberg (32 G6A)**
- **Cambridge (1 G6)**
- **Edinburgh (1 G6)**
- **Amsterdam (4 G6)**
- **Bonn**
- **Munich**

Science with GRAPE

- Solar system/Planetary formation Daisaka, Takeda
- Star-forming region/Open clusters
- Globular clusters
- Galactic nuclei
- Galaxy formation Saitoh, Nakasato
- Clusters of Galaxies/Cosmology

4 out of 6 talks on N-body simulation this afternoon discuss GRAPE results. (excluding this talk)

Massive Black Hole in GCs?

Observation + Interpretation

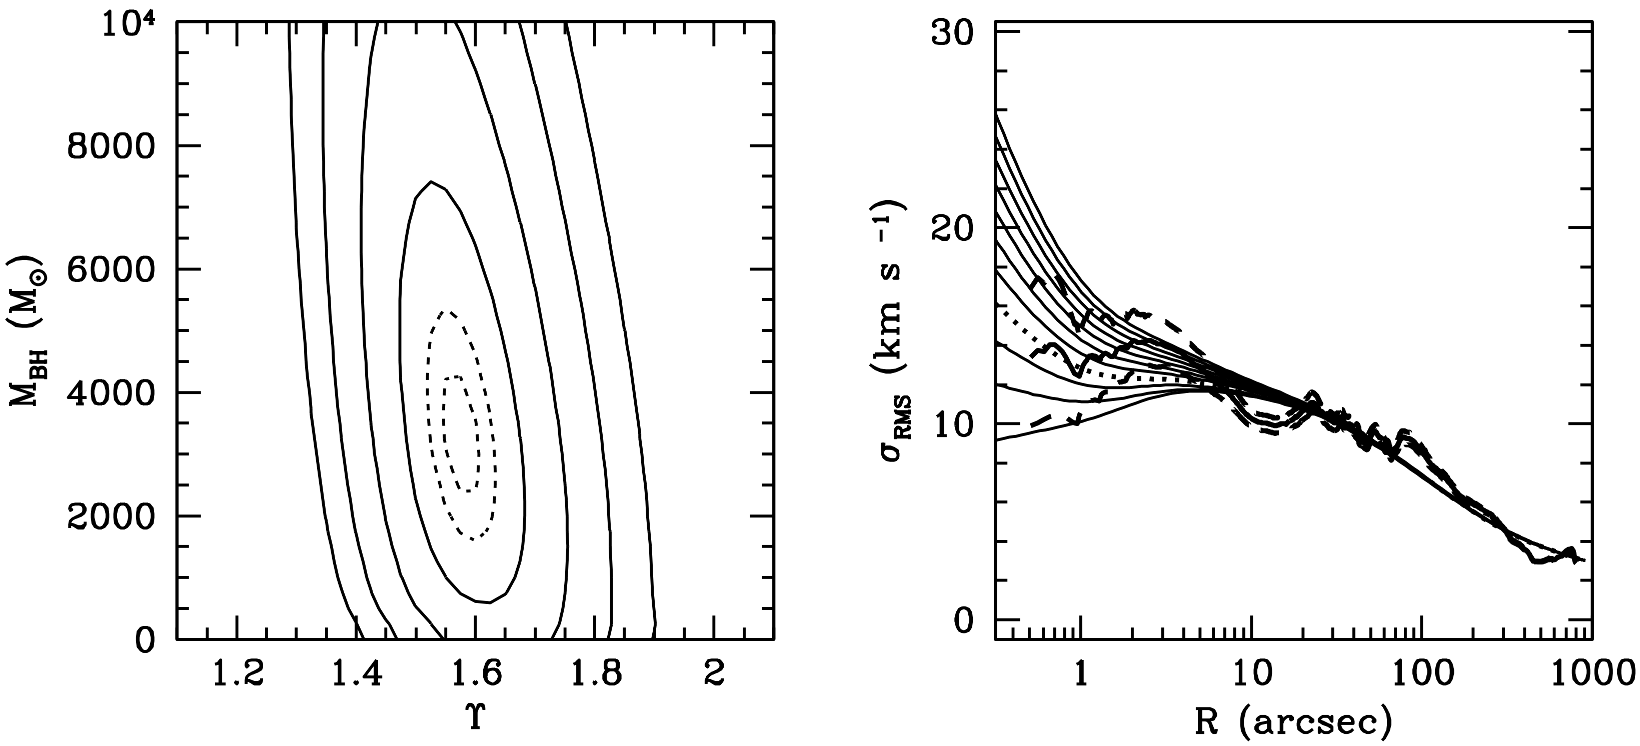

3000 $M_\odot$ black hole in M15? (Gerssen et al 2002)
Without BH, luminocity profile gives too low central velocity dispersion.

N-body simulation without BH

Baumgardt et al., ApJ 2003, 582, L21.

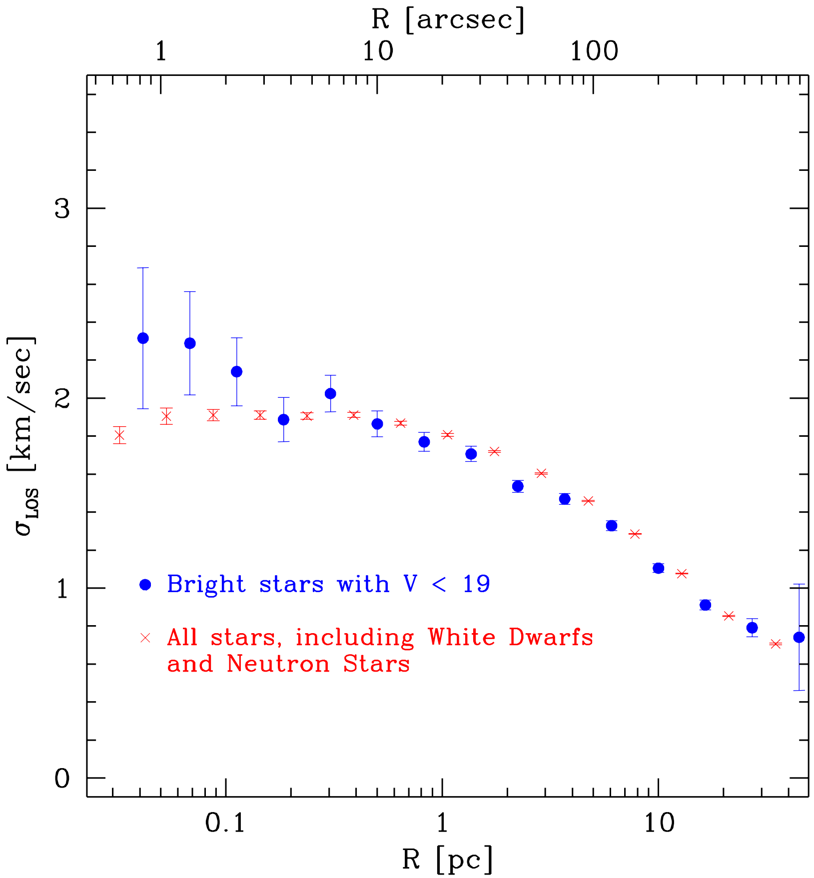

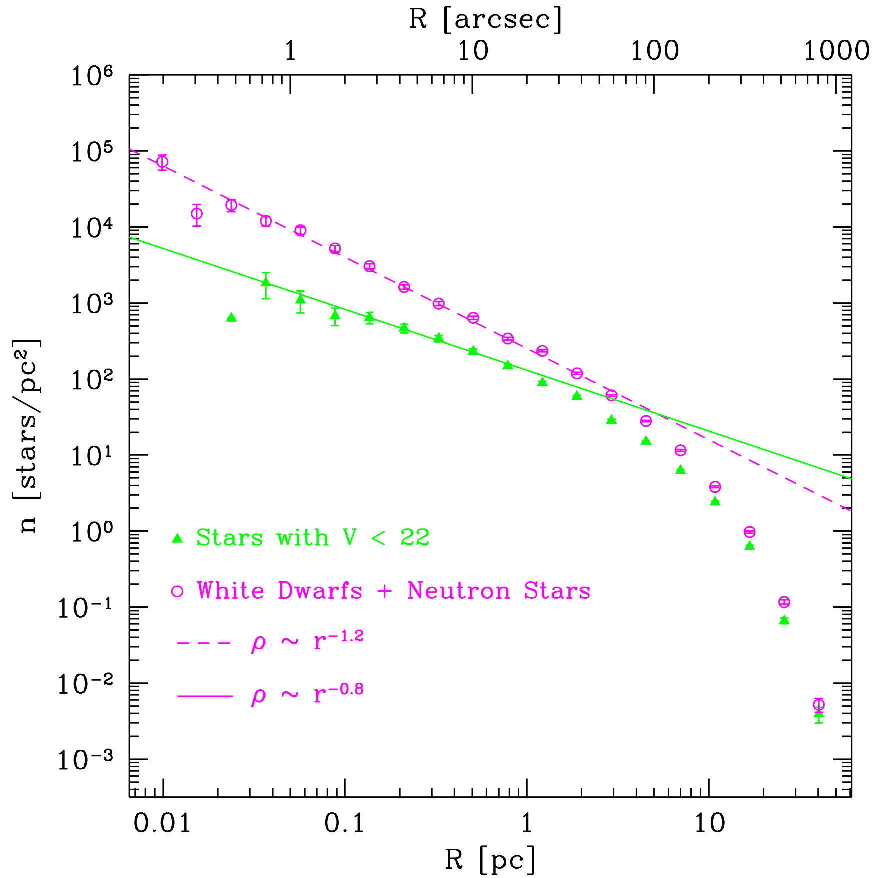

Left: velocity dispersion; Right: Surface density.

We "found" BH, though there wasn't

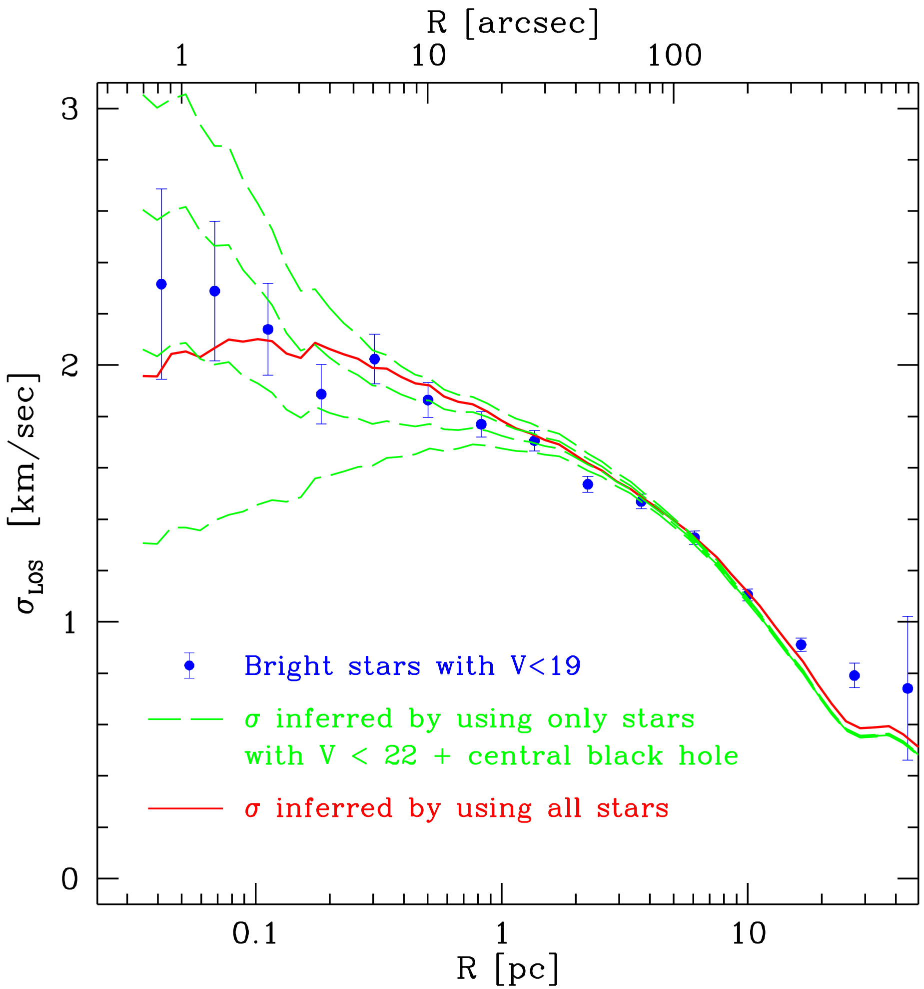

By analyzing the simulation data in the same way as the observers did, we "found" a central black hole, though it is not there.

Estimated BH mass = $80M_\odot$. If scaled to M15, $\sim 3 \times 10^3 M_\odot$ (Gerssen *et al.*: $\sim 3 \times 10^3 M_\odot$)

M 15 does not need black hole.

Is there any GC with central BH?

Baumgardt et al. 2004 (ApJL, submitted)

How would a GC with central BH look like?

$$\downarrow$$

Simulate globular clusters with central BH for a Hubble time.

Profile evolution

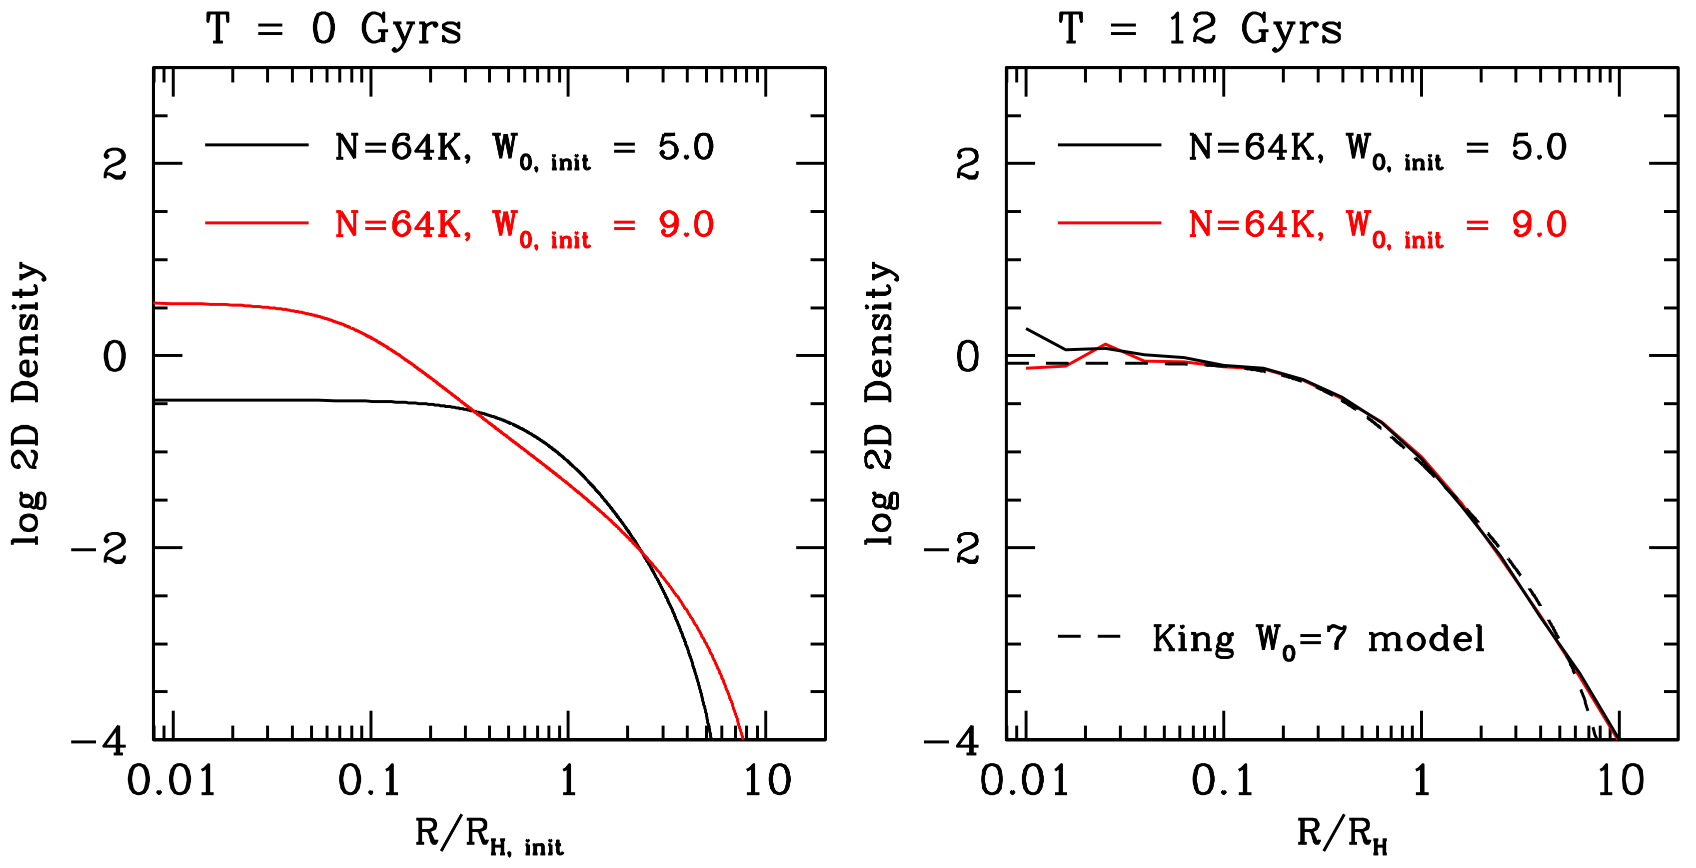

- Surface brightness profile becomes King7-like.
- Almost independent of initial profile and BH mass (in the range of 0.1% to 1%)

Globular cluster summary

- Globular clusters with central luminosity cusp do not contain massive central BH. They are really clusters in deep core collapse, with NS and WD dominating the central cusp.
- Most likely place to find massive central BH is some of normal-looking clusters with relatively large cores.

Next-Generation GRAPE — GRAPE-DR

- Budget approved. (1.5M$ $\times$ 5 years)
- Planned peak speed: 2 Pflops
- New architecture — wider application range than previous GRAPEs
- Planned completion year: 2008

Why new architecture?

Essential reason:

Economical/Political, not technical/scientific...

Custom chip costs too much

1990	1μm	150K$
1997	0.25μm	1M
2004	90nm	> 3M

- Theoretical Astrophysics in Japan is too narrow.
- We need to widen the application range.
- Wider application range (if possible) also helps in astrophysics.
  - SPH
  - Radiative transfer etc...

GRAPE-DR design goal

- **Keep the high performance of previous GRAPE architecture**
- **Make it "programmable", at least in some limited sense**

GRAPE-DR design goal

- **Keep the high performance of previous GRAPE architecture**
- **Make it “programmable”, at least in some limited sense**

How?

GRAPE-DR design goal

- **Keep the high performance of previous GRAPE architecture**
- **Make it “programmable”, at least in some limited sense**

How?

Key: High degree of on-chip parallelism

Parallelism in GRAPE

Formula to evaluate:

$$\boldsymbol{a}_i = \sum_j \boldsymbol{f}(\boldsymbol{r}_i, \boldsymbol{r}_j, \boldsymbol{m}_j)$$

Parallelism exists for both i and j

GRAPE uses three hardware-level parallelism:

Pipelined processor	**j-parallelism**
Multiple pipeline in a chip	**i-parallelism**
Multiple chips	**both way**

Most of parallelism is realized as multiple processors.
Pipelined processor is only a small part of the solution.

GRAPE-DR Architecture

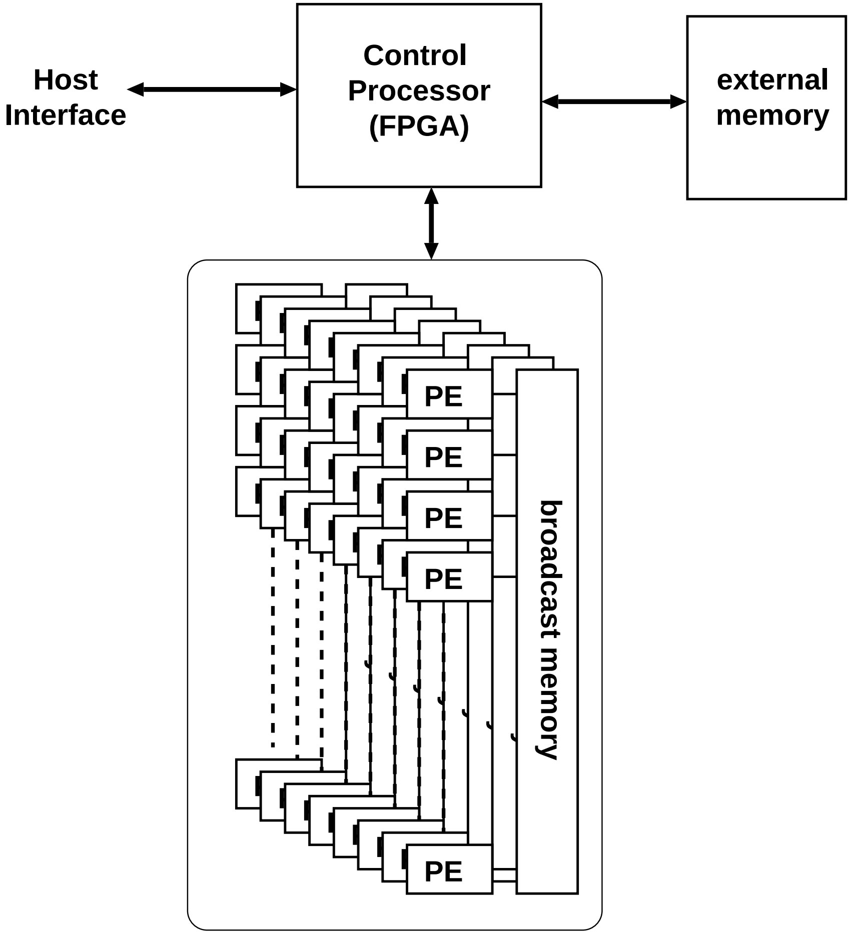

- Eliminate hardwired pipeline
- Integrate huge number of very simple processors ($\sim 1,000$) to a chip
- processors (PE) = ALU+register file (no local memory)
- ALL PE run in parallel under the control of a single program

Parallel processor and pipeline

- Anything a pipeline can do a parallel processor also can do.

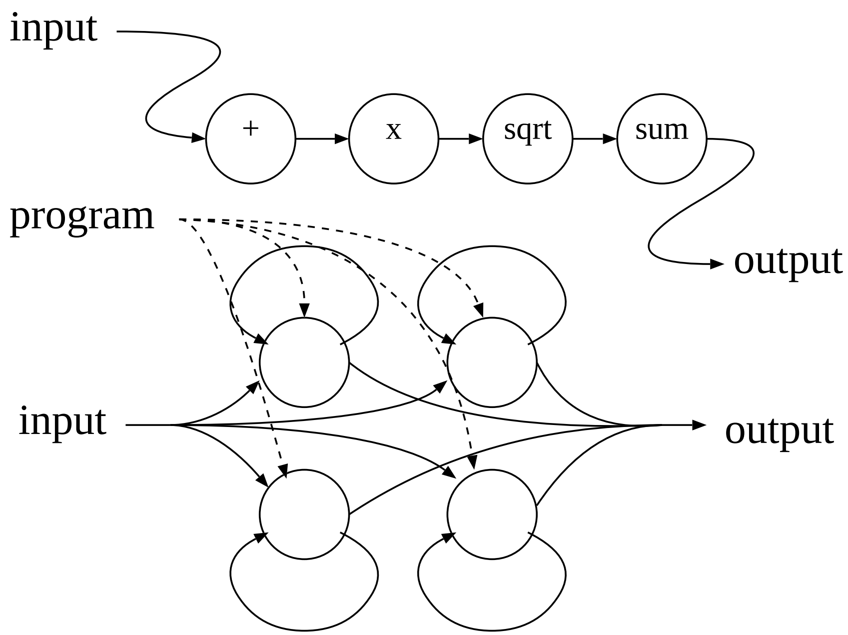

- Specialized pipeline processor for a specific application can achieve better performance.

- This advantage is, however, less than a factor of 10 and does not depend on the semiconductor technology.

PE Structure

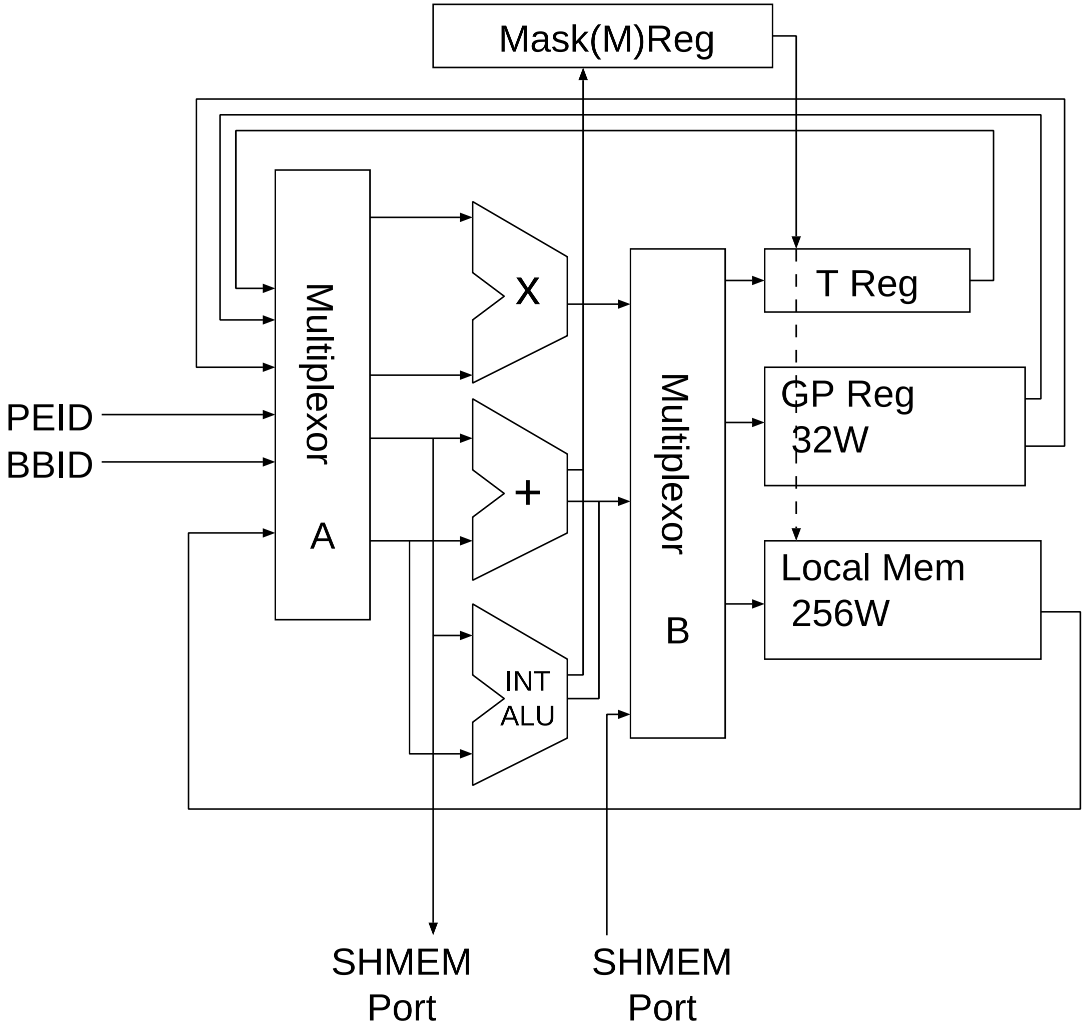

- Floating-point unit
- Integer ALU
- registers
- memory(2KB)

GRAPE-DR processor structure

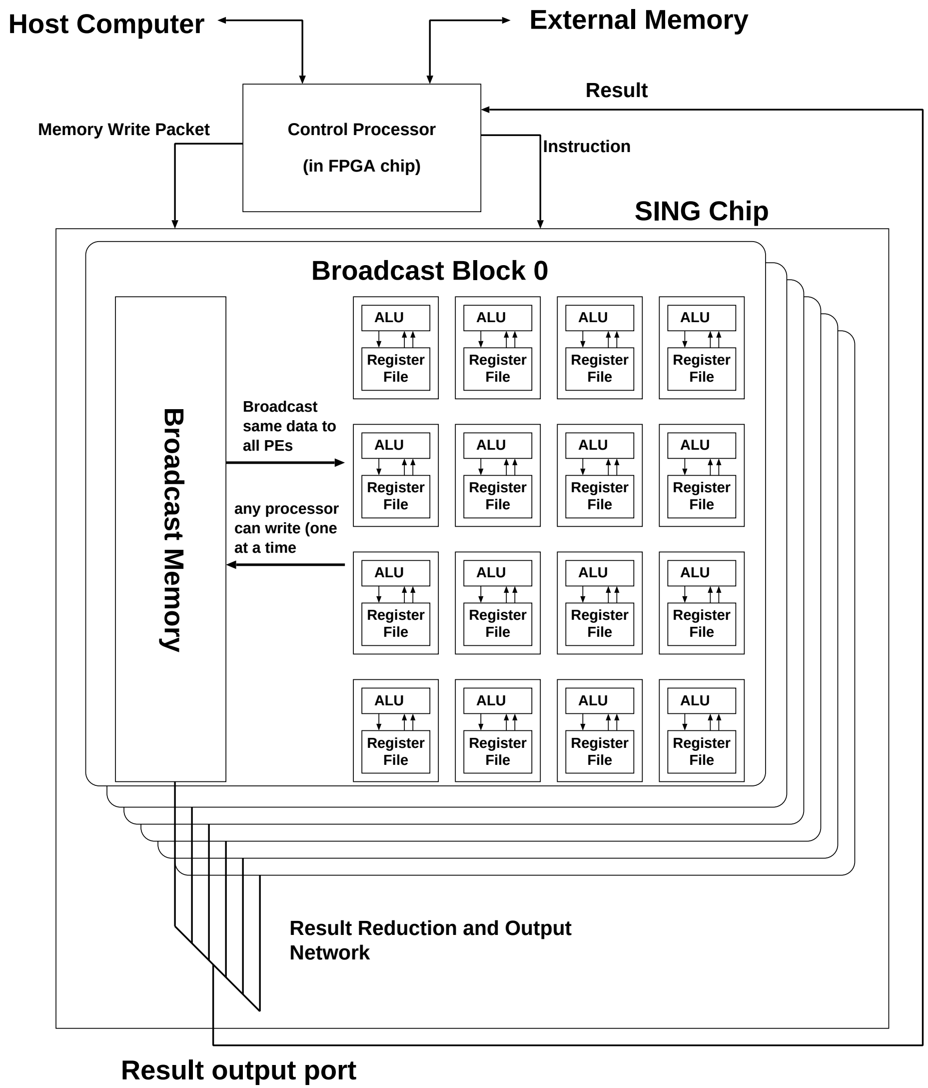

Collection of small processor, each with ALU, register file (local memory)

One chip will integrate (hopefully) 1024 processors
Single processor will run at 700MHz clock (2 operations/cycle).

Peak speed of one chip: 1-2 Tflops (30-60 times faster than GRAPE-6).

GRAPE-DR system structure

- **4 chips per board (PCI-X or PCI-Express)**
- **2 boards per node (a node is standard PC with 10GbE network)**
- **512-node cluster system**

Total peak speed: 2-3 Pflops.

Other applications

- Particle-based simulations like MD, SPH
- Dense matrix operation(Linpack, LU decomposition, Eigenvalue calculation)
- BEM: Poisson equation, Electromagnetics
- CFD with spectral method over sphere
- Molecular orbital calculation

Applications that need

- large amount of computation
- but not much memory or memory bandwidth

are potential targets.

Summary

- GRAPE project has successfully developed very high performance computers for astrophysical particle-based simulations.
- The next machine, GRAPE-DR, will have wider application range than traditional GRAPEs, keeping good price-performance ratio.

Emulating GRAPE

In the simplest case:

- **Each PE calculates the force on its particle, but from one same particle**
- **The particle which exerts the force is supplied from the external memory.**