

Japanese Supercomputer Project

Jun Makino

University of Tokyo

Japanese Supercomputer Project

is the title of my talk in the program

Japanese Supercomputer Project

is the title of my talk in the program

However, I do not know much about Japanese projects in general (apart from the names, Earth Simulator, PACS-CS, Post Earth Simulator)

Japanese Supercomputer Project

is the title of my talk in the program

However, I do not know much about Japanese projects in general (apar from the names, Earth Simulator, PACS-CS, Post Earth Simulator)

So I will talk about our project.

Current status of GRAPE Project

Jun Makino

University of Tokyo

Talk overview

- **GRAPE Project**
- **GRAPE-DR: Next-Generation GRAPE**

GRAPE project

- basic idea
- hardware
- usage example
- GRAPEs in the world

GRAPE project: Rationale

GOAL:

Design and build specialized hardware for simulation of stellar systems.

Rational:


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\$
BG/L	(2005?, 360 TF?)	??M\$

Basic idea of GRAPE

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

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



Time integration,
IO, etc

Force calculation

Flexibility

High performance

GRAPE machines

1989	GRAPE-1	240 MF	Low accuracy(LA)
1990	GRAPE-2	40 MF	High accuracy(HA)
1991	GRAPE-3	15 GF	LA, custom chip
1995	GRAPE-4	1.08 TF	HA, custom chip
1998	GRAPE-5	40*n GF	LA, 2 pipelines in a chip
2001	GRAPE-6	64 TF	HA, 6 pipelines in a chip


Molecular Dynamics

1992	GRAPE-2A	120MF	
1996	MD-GRAPE	2.4GF	custom chip
2001	MDG2(MDM)	75 TF	RIKEN
2006?	MDG3	0.6PF	RIKEN



Next Generation

2008?	GRAPE-DR	2PF	
2009?	MDG4?	16PF?	
2001x?	GDR2??	16PF???	

Evolution of peak performance



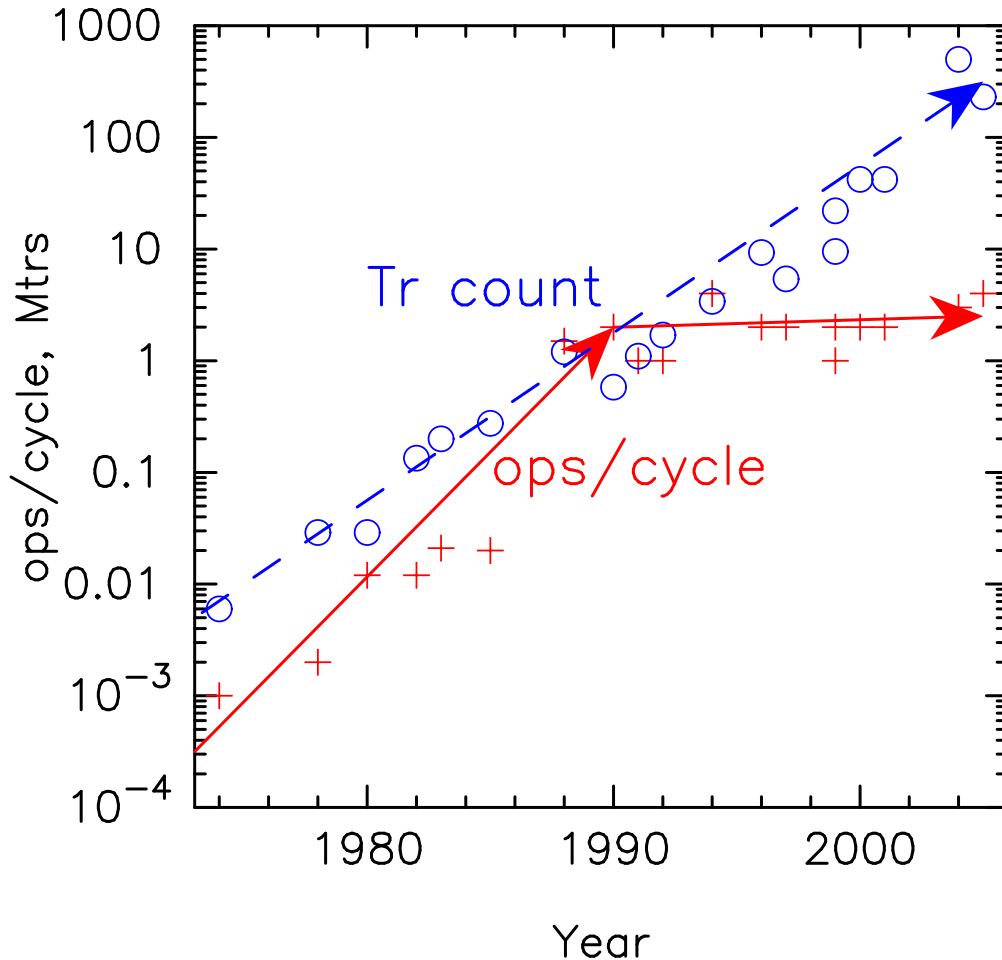
Intel P4 and GRAPE-6



Intel Prescott (2004)
2 FP ops/clock
90nm, 7.6 GF, > 100W?

GRAPE-6 chip (2000)
~ 400 FP ops/clock
0.25μm, 31 GF, 10W

Evolution of microprocessors




Number of transistors doubles every 18 months (*“Moore’s Law”*)

Number of floating point units got stuck at $O(1)$. Never reached more than 4.

You can do much better than COTS microprocessor if you can use more than 10% of transistors for FP operations!


The GRAPE approach

General-purpose



How to connect processors and memories???

GRAPE




Hardwired pipeline
One memory serves many pipelines

GRAPE-6


- processor chip
- processor board
- total system

Pipeline LSI

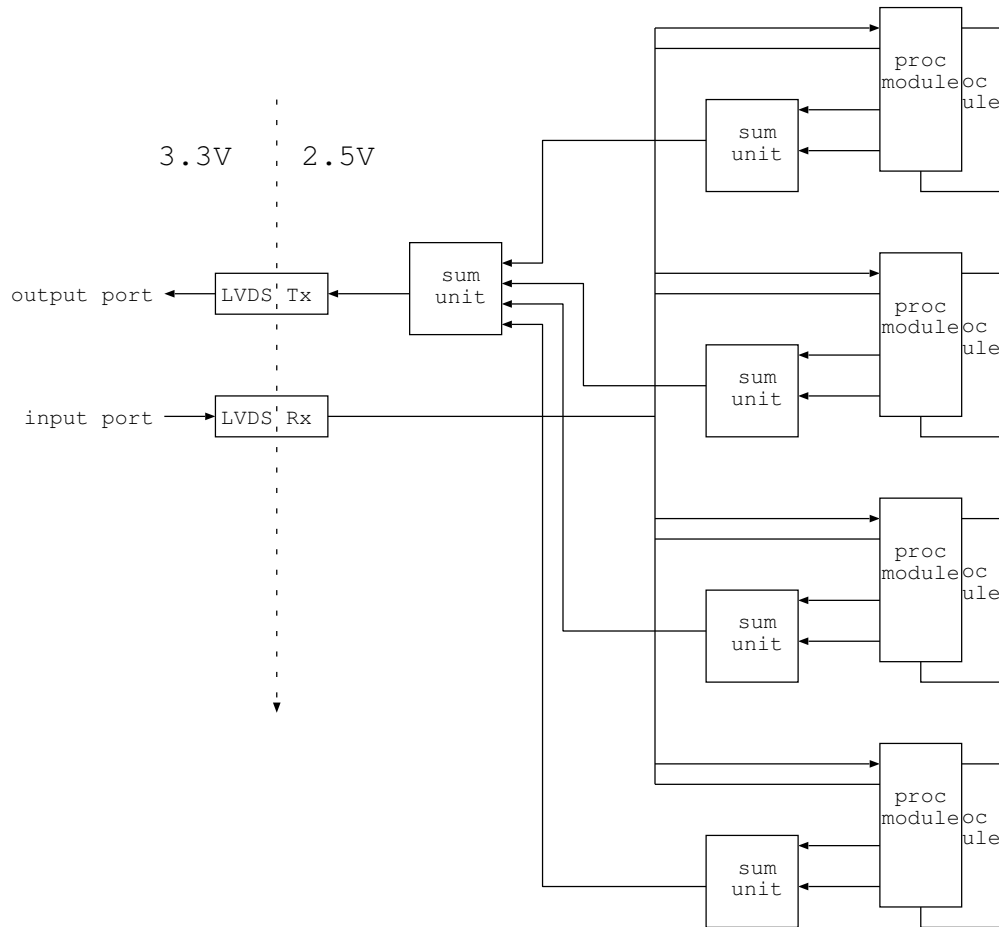


- 0.25 μm design rule
(Toshiba TC-240,
1.8M gates)
- 90 MHz clock
- 6 pipelines
- one predictor pipeline
- 31 Gflops /chip

GRAPE-6 processor board




GRAPE-6 Processor board



- 32 chips/board
- Semi-serial (LVDS) interface (350MHz clock, 4 wires)
- Tree network for data broadcast and reduction

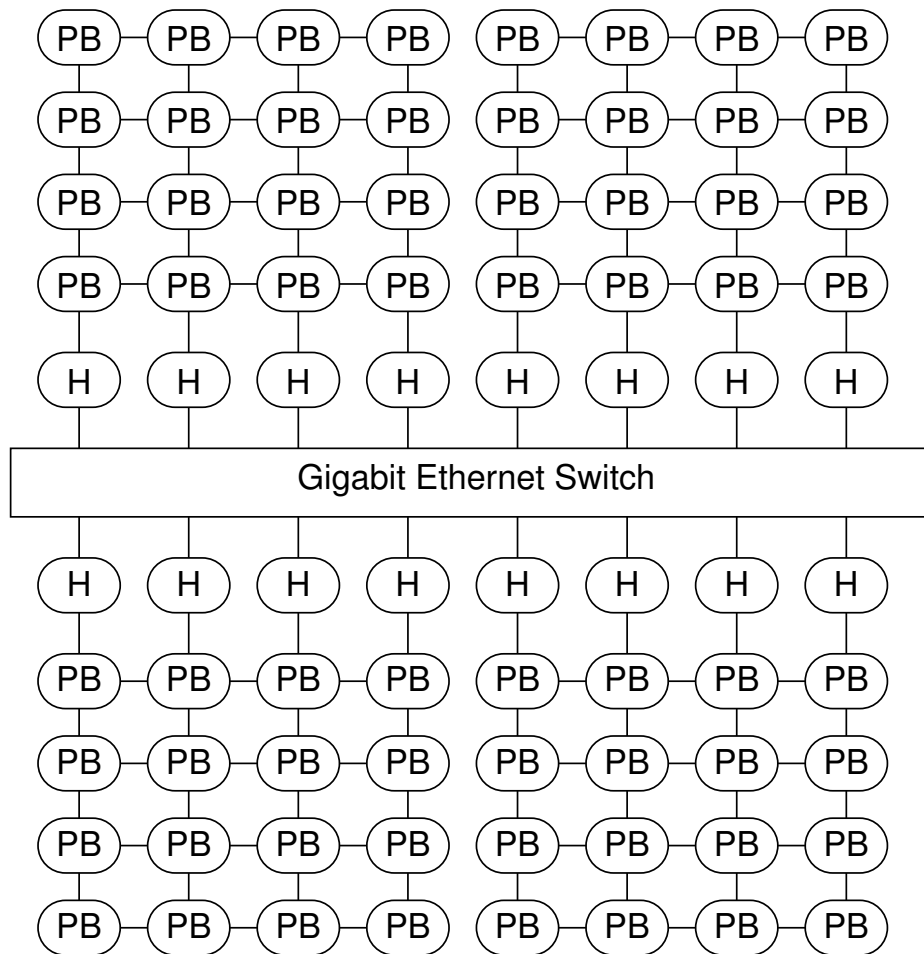
The 64-Tflops GRAPE-6 system



Present 64-Tflops system.

4 blocks with 16 host computers.

The full 64 Tflops GRAPE-6 system



- 4-host, 16-board “block” with dedicated network
- 4 “blocks” connected through GbE network


Combination of host network solution and dedicated network solution.

Some performance numbers

Gordon Bell Prizes

- 2000 1.34 TF on 64-chip system
- 2001 11.2 TF on 1024-chip system
- (2002 29.5TF on 2048-chip system)
- 2003 **33.4TF** 2048-chip system

BabyGRAPE (aka microGRAPE)



Fukushige et al 2005

Single PCI card with peak speed of 123 Gflops

Commercial version: http://www.metrix.co.jp/micro_grape_eng.html

GRAPE6 worldwide

incomplete list of GRAPE-6s

AMNH 4 G6s

Amsterdam

ARI Heidelberg **32** BGs

Bonn

Cambridge

Drexel 2 G6s?

Indiana

Marseilles

McMaster

Michigan

MPIA

Munich

NAOJ 12 G6s

Rochester **32** BGs

TIT

Tsukuba **256** BGs (06?)


Science with GRAPE

- Cosmology (CDM halo)
- Globular clusters
- Galactic nuclei (black hole binaries)
- Planet formation
- Star formation
- Young star cluster (Portegies Zwart)
- Galactic dynamics
- galaxy formation
- ...
- ...

CDM halo simulation




GRAPE-5 Cluster
8 nodes, 40 Gflops each



Simulated Cluster
60M particles in a single DM
halo
Largest calculation of this kind.

Globular clusters with central black hole

Baumgardt, J.M. and Hut (ApJL 620, 238, 2005)

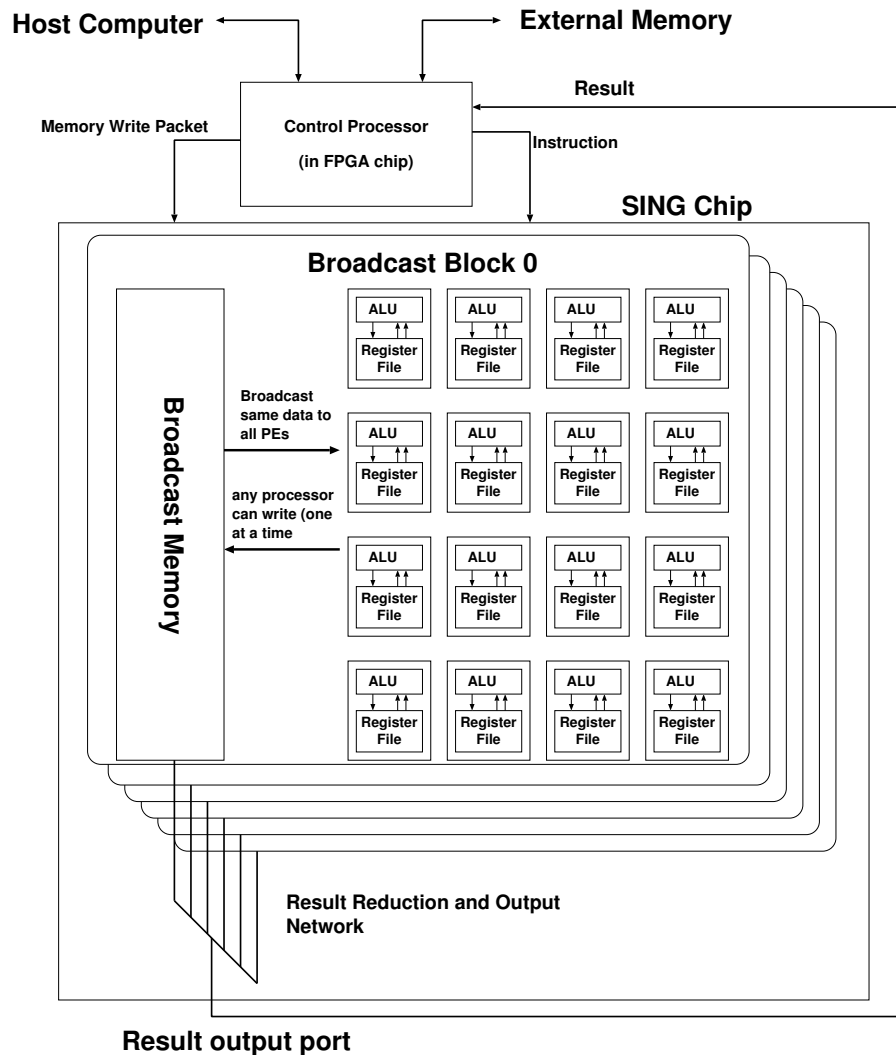


Surface brightness profile becomes King7-like, almost independent of initial profile and BH mass (in the range of 0.1% to 1%)

Next-Generation GRAPE — GRAPE-DR

- Budget approved. (2.5M\$ × 5 years)
- Planned peak speed: 2 Pflops
- New architecture — wider application range than previous GRAPEs
- Planned completion year: FY 2008 (early 2009)

GRAPE-DR processor structure



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

One chip will integrate 512 processors

Single processor will run at 500MHz clock (2 operations/cycle).

Peak speed of one chip: 0.5 Tflops (20 times faster than GRAPE-6).

High-level architecture

- Single card: 4 chips, PCI-X/PCI-E/Hypertransport(? interface, 2 Tflops.
- Two cards per single node, host: x86 PCs.
- Host network: 512 node, fast 10GbE switches?

Difference from GRAPE-6:

- No custom network
- No large card

Development schedule

(tentative)

2006 Fall	First sample chip
2007 Spring	Prototype board
2008 Spring	Large parallel system
2009 Spring	Final system

Summary

- GRAPE project has successfully developed very high performance computers for astrophysical particle-based simulations.
- This high performance is achieved by designing a custom processor chip with very large number of arithmetic units, connected in the form of application-specific hardwired pipeline.
- The next machine, GRAPE-DR, will have wider application range than traditional GRAPEs, with programmable processors.