

\$105/Gflops

Astrophysical  $N$ -body Simulation with  
Reconfigurable Add-in Card and  
Hierarchical Tree Algorithm

A. Kawai\* & T. Fukushima\*\*

\* Saitama Institute of Technology

\*\* University of Tokyo, K&F Computing Research Co.

# Abstract

Reconfigurable  
Hardware

+

Tree  
Algorithm

The first Gordon-Bell finalist.




Achieved Price/Performance: \$105/Gflops

- Less than half of the current record.
- 67 times better than the same run in GB99.


# Reconfigurable Hardware

FPGA based add-in card.

PCI-sized  
add-in card

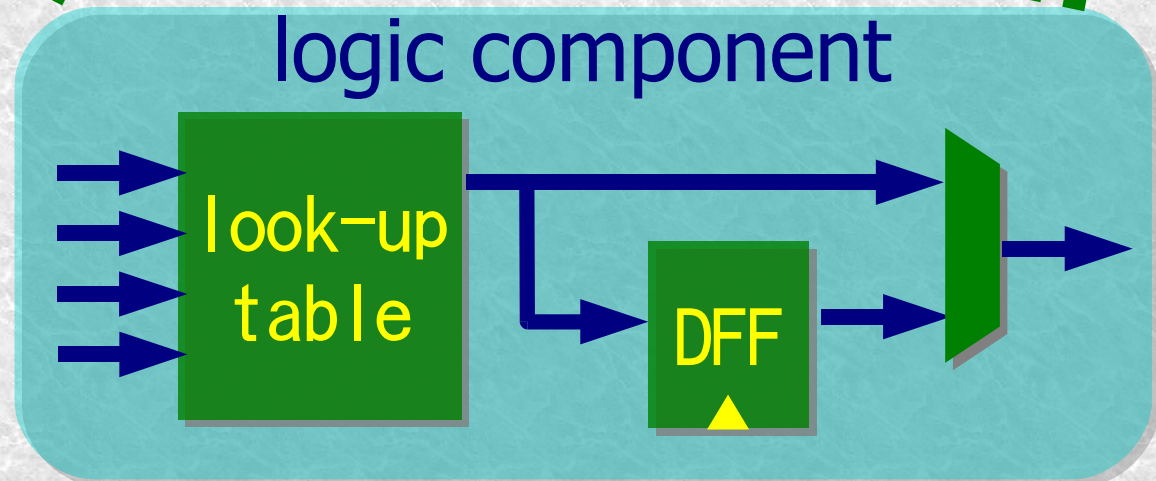


FPGA



array of  
programmable  
logic

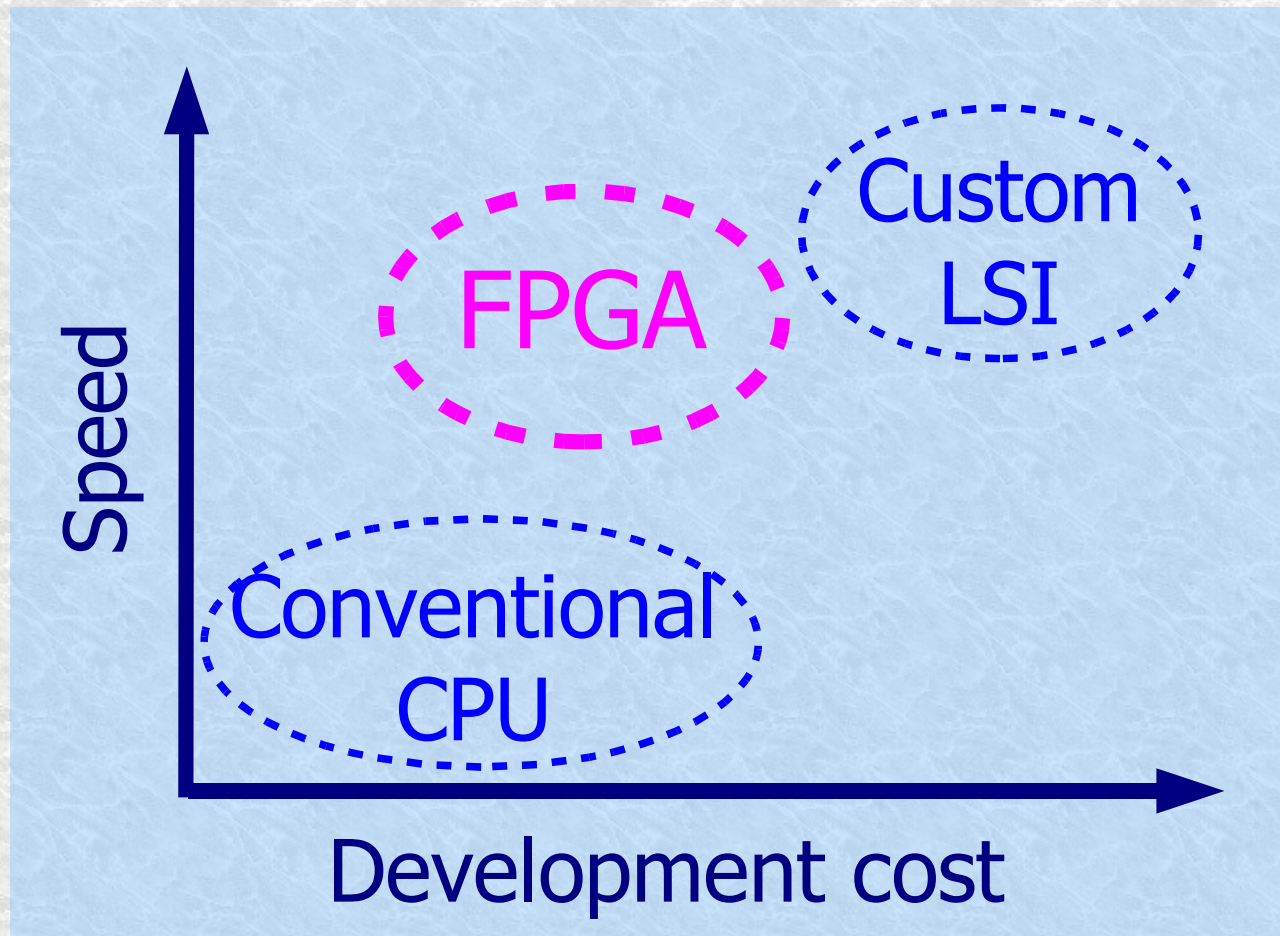
combinational  
logic by  
look-up table



# Advantage of Reconfigurable Hardware

Offers reasonable balance of speed & cost, if :

- a huge # of reproduction is not necessary.
- applications require modest accuracy.





# Reconfigurable Hardwares on the Market

## FPGA Vendors

Altera: Stratix series, Cyclone series, ...

Xilinx: Virtex series, Spartan series, ...

Lattice: LatticeSC, LatticeEC, ...

## Supercomputers Equiped with FPGAs

Cray XD1: Xilinx Virtex-4 + Opteron

SRC MAPstation: Xilinx Virtex-4 + Xeon

SGI RASC: Xilinx Virtex-4 (+ host computer)

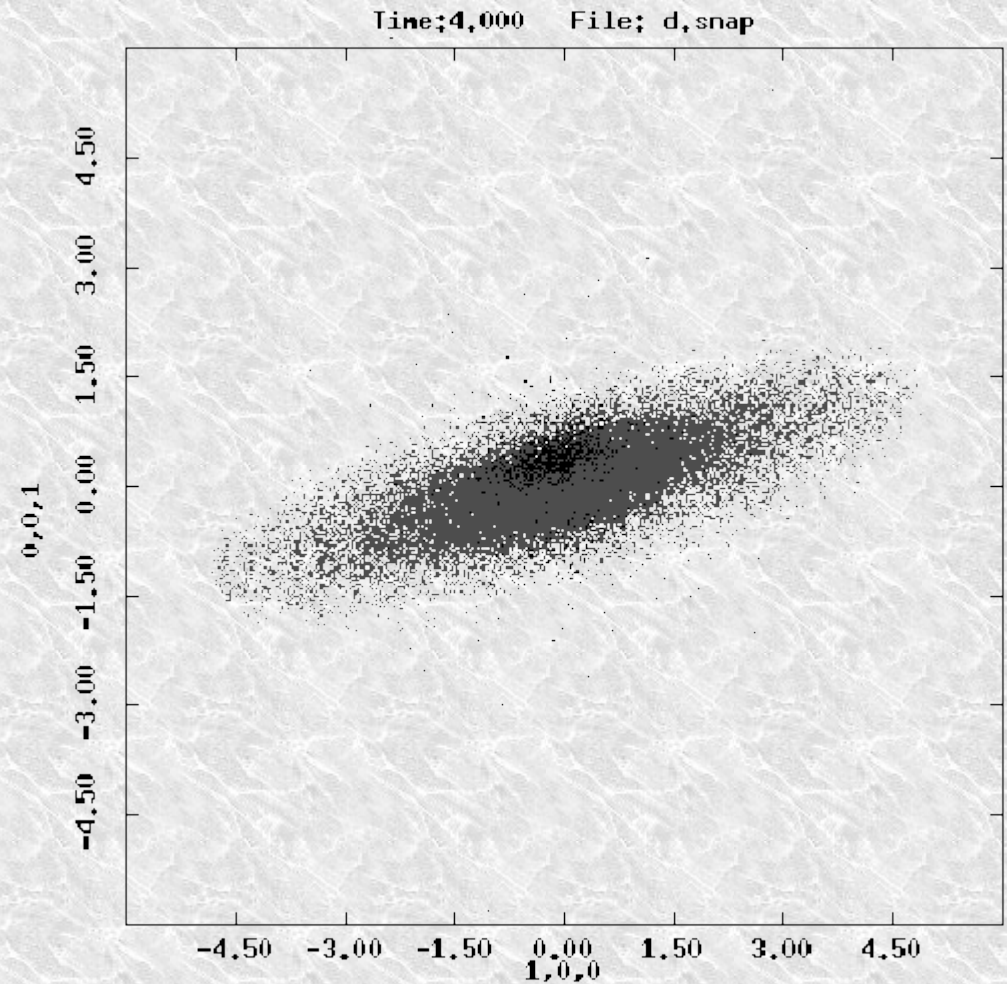
# Astrophysical Simulation in HPC

- Astrophysical  $N$ -Body Simulation
- Calculation Cost
- Past Winners' Approach

# Astrophysical $N$ -Body Simulation (1)



An astronomical object




A representation by particles

# Astrophysical $N$ -Body Simulation (2)

Equation of Motion:


$$\begin{cases} \frac{\partial^2 \vec{r}_i}{\partial t^2} = \vec{f}_i \\ \vec{f}_i = \sum_{j \neq i}^N G \frac{m_j (\vec{r}_j - \vec{r}_i)}{|\vec{r}_j - \vec{r}_i|^3} \end{cases}$$






# Calculation Cost of $N$ -Body Simulation

## Direct Summation Algorithm



$O(N^2)$  : Prohibitive

## Tree Algorithm




$O(N \log N)$  : Large

➡ A "Grand Challenge" problem in HPC.

# Astrophysical Simulations in GB History

Performance  
(Tflops)

Price/Performance  
(\$/Mflops)



## Past Approaches

Massively-Parallel System + Tree

Warren et al.(1998,1997),cosmology

Dedicated Hardware + Tree

Kawai et al.(1999),cosmology

Dedicated Hardware +

Direct sum. + multiple timestep

Makino et al.(2003),proto-planet system

Makino et al.(2000,2001),black hole system

## Our Approach

Reconfigurable Hardware + Tree


Kawai et al.(2006),cosmology

# Detail of Our Simulation

- Hardware Configuration
- Simulation
- Price/Performance
  - Performance
  - Performance Correction
  - Price/Performance



# Hardware Configuration




## Tested Systems:

Core 2 Duo E6400	+	KFCR GRAPE-7 model100
PentiumD 920	+	PLDA PXPDKSXGX40
Pentium4 630	+	PLDA XpressBridge






KFCR GRAPE-7  
model100



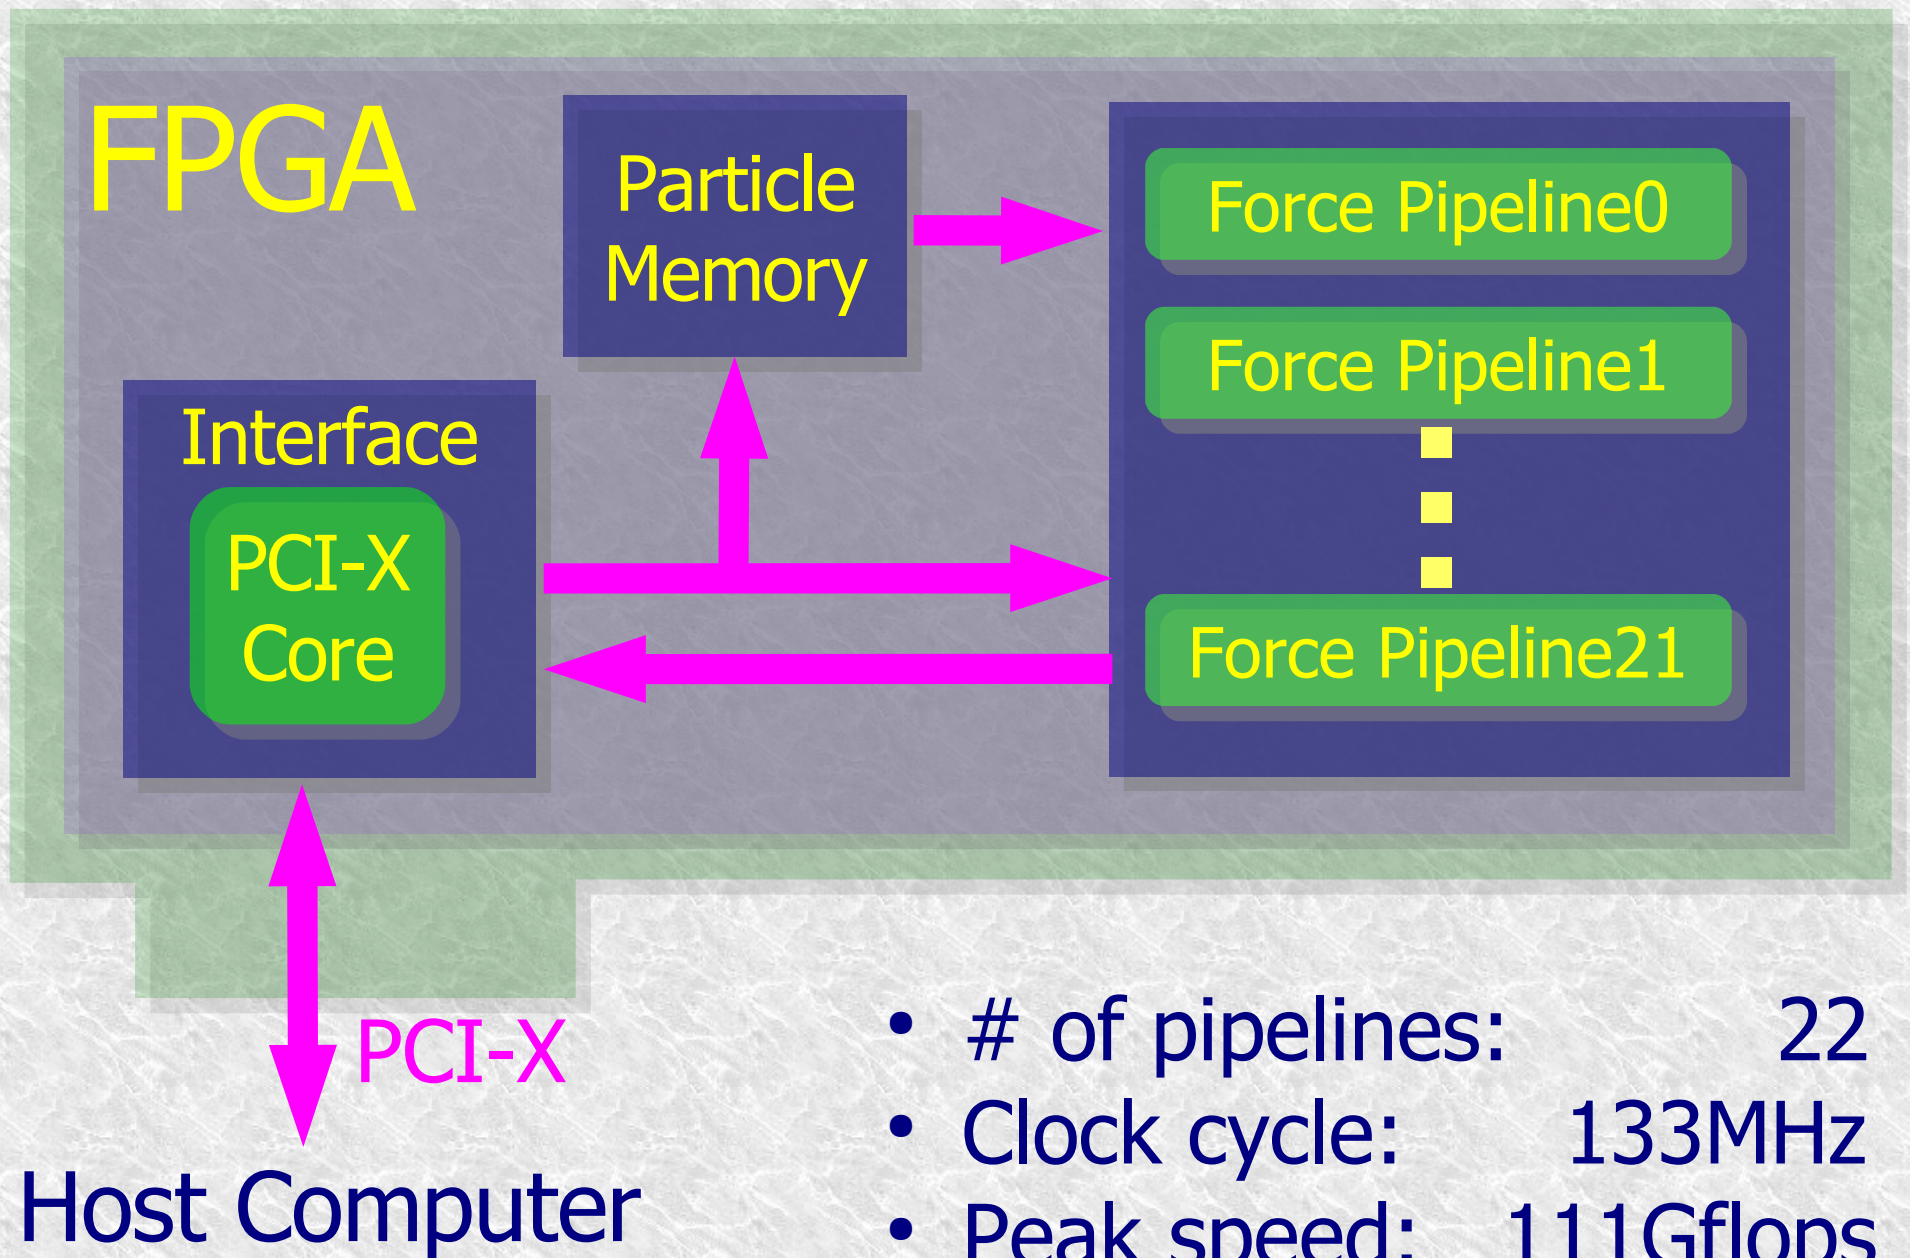
PLDA PXPDKSXGX40



PLDA XpressBridge




# Add-in Card Configuration



- # of pipelines: 22
- Clock cycle: 133MHz
- Peak speed: 111Gflops
- Power consumption: 3W

# Modified Tree Algorithm (Barnes1990)

Modification for efficient parallelization.




One interaction list is **shared** among  $n_g$  particles.

	Original	Modified
List creation on the host:	$N$ times	$N/n_g$ times
Interaction list length:	shorter	longer



# A Cosmological Astrophysical Simulation



$N : 2.16 \times 10^6$

$\theta : 0.4$

timestep : 1000 steps

$\rho$  : based on a standard  
cold dark matter  
scenerio.

# Performance

Elapsed time: 2.95 hours

Operation per interaction: 38 flop

Interaction count:  $1.49 \times 10^{13}$

Performance: 53.16 Gflops (48% of peak)

correction



Interaction count:  $6.31 \times 10^{12}$

Performance: 22.59 Gflops

# Price/Performance

Price (JPY)

Add-in Card: 187,619

Host PC: 91,238

---

Total: 278,857



\$2,363

\$1=118JPY

Corrected Performance: 22.59 Gflops

Price/Performance: \$105/Gflops

(improved from \$158/Gflops marked in July)



Previous record: \$246/Gflops (Kim et al. 2001)




# Comparison with Our Previous Run

	2006 Reconfigurable System	1999 GRAPE-5 System
Technology	FPGA(90nm)	ASIC(500nm)
# of pipelines	22	2/chip x 16 chips
Clock cycle	133MHz	90MHz
Dimensions	short PCI	39x31x7(cm) x 2chasis
Power consumption	3W	70W x 2
Interface	PCI-X	PCI
Host CPU	Core 2 duo	Alpha21264
Speed (sustained)	23Gflops	6Gflops
Cost	\$2,363	\$40,900
Price/perf.	\$105/Gflops	\$7000/Gflops




# Scalability

Preliminary test on a small cluster (4cards+2hosts)



Price/perf. achieved with 1-card: \$158/Gflops  
4-card: \$202/Gflops



Expected to show good scalability when fully tuned.

# Summary

Reconfigurable Hardware

force calculation  
pipelines

+

Tree Algorithm



Price/Performance: \$105/Gflops

- Less than half of the Gordon-Bell record.
- 67 times better than the same run in GB99.