

GRAPE-6: Massively-Parallel Special-Purpose Computer for Astrophysical Particle Simulations

Junichiro MAKINO

*Department of Astronomy, School of Science, The University of Tokyo, Tokyo 133-0033
makino@astron.s.u-tokyo.ac.jp*

Toshiyuki FUKUSHIGE and Masaki KOGA

*Department of General System Studies, College of Arts and Sciences, The University of Tokyo, Tokyo 153-8902
fukushig@provence.c.u-tokyo.ac.jp*

and

Ken NAMURA

*IBM Japan Industrial Solution Co., Ltd, Yamato, Kanagawa 242-8502
JL26165@jp.ibm.com*

(Received 2003 June 10; accepted 2003 October 8)

Abstract

In this paper, we describe the architecture and performance of the GRAPE-6 system, a massively-parallel special-purpose computer for astrophysical N -body simulations. GRAPE-6 is the successor of GRAPE-4, which was completed in 1995 and achieved the theoretical peak speed of 1.08 Tflops. As was the case with GRAPE-4, the primary application of GRAPE-6 is simulations of collisional systems, though it can also be used for collisionless systems. The main differences between GRAPE-4 and GRAPE-6 are (a) the processor chip of GRAPE-6 integrates 6 force-calculation pipelines, compared to one pipeline of GRAPE-4 (which needed 3 clock cycles to calculate one interaction), (b) the clock speed is increased from 32 to 90 MHz, and (c) the total number of processor chips is increased from 1728 to 2048. These improvements resulted in a peak speed of 64 Tflops. We also discuss the design of the successor of GRAPE-6.

Key words: celestial mechanics — methods: N -body simulations

1. Introduction

The N -body simulation technique, in which the equations of motion of N particles are integrated numerically, has been one of the most powerful tools for studying astronomical objects, such as the solar system, star clusters, galaxies, clusters of galaxies, and large-scale structures of the universe.

Roughly speaking, the target systems for N -body simulations can be classified into two categories: collisional systems and collisionless systems. In the case of collisional systems, the evolution of the system is driven by a two-body relaxation process, in other words, by the microscopic exchange of thermal energies between particles. In this case, the simulation timescale tends to be long, since the relaxation timescale measured by the dynamical timescale is proportional to $N/\log N$, where N is the number of particles in the system.

The calculation cost of the simulation of collisional systems increases rapidly as we increase the number of particles, N , for the following two reasons. First, as stated above, the relaxation timescale increases roughly linearly as we increase N . This means that the number of timesteps also increases at least linearly (Makino, Hut 1988). The second reason is that it is not easy to use fast and approximate algorithms, such as Barnes–Hut tree algorithm (Barnes, Hut 1986) or the fast multipole method (Greengard, Rokhlin 1987), to calculate the interaction between particles. These facts imply that the cost per timestep is $O(N^2)$, and that the total cost of the simulation is $O(N^3)$.

There are two reasons why the use of approximate

algorithms for the force calculation is difficult. The first reason is the need for relatively high accuracy. Since the total number of timesteps is very large, we need a rather high accuracy for the force calculation. The other reason is the wide difference in the orbital timescale of particles. A unique nature of the gravitational N -body problem is that particles interact only through gravity, which is an attractive force. This means that two particles can approach arbitrary close during a hyperbolic close encounter. In addition, spatial inhomogeneity tends to develop, resulting in a high-density core and a low-density halo. Even on average, particles in the core require much smaller timesteps than do particles in the halo.

It is clearly very wasteful to apply the same timestep to all particles in the system, and it is crucial to be able to apply an individual and adaptive timestep to each particle. Such an “individual timestep” algorithm, first developed by Aarseth (1963, 1999), has been the core for practically any program that handles the time integration of collisional N -body systems, such as star clusters and systems of planetesimals.

The basic idea of the individual timestep algorithm is to assign different times and timesteps to particles in the system. For particle i , its next time is $t_i + \Delta t_i$, where t_i is the current time and Δt_i is the current timestep. To integrate the system, we first choose a particle with minimum $t_i + \Delta t_i$ and set the current system time t to be $t_i + \Delta t_i$. Then, we predict the positions of all particles at time t and calculate the force on particle i . Finally, we correct the position of particle i using the calculated force, update t_i and determine the new timestep

Δt_i . In practice, we force the size of timesteps to be powers of two, so that the system time is quantized and multiple particles have exactly the same time. In this way, we can use parallel or vector processors efficiently, since we can integrate multiple particles in parallel (McMillan 1986; Makino 1991a).

It is necessary to use the linear multistep method (predictor–corrector method) with variable stepsize for the time integration. Aarseth adopted an algorithm with third-order Newton interpolation. Recently, the method based on the third-order Hermite interpolation (Makino 1991b; Makino, Aarseth 1992) has become widely used, because of its simplicity.

In principle, it is not impossible to combine the individual timestep algorithm and fast algorithms, such as the Barnes–Hut tree algorithm or FMM. McMillan and Aarseth (1993) developed such a combination, where the tree structure is dynamically updated according to the motion of the particles, and force is calculated using multipole expansion up to octupole. They assigned predictor polynomials to each node of the tree structure so that they could calculate the force from nodes to particles at arbitrary times.

A serious problem with such a combination is that there is no known method to implement it on parallel computers with distributed memory. It is not simple to achieve a good parallel performance with the individual timestep algorithm, even without the tree algorithm. The reason is that simple methods require fast and low-latency communication between processors. A recently proposed two-dimensional algorithm (Makino 2002) somewhat relaxes the requirement for the communication bandwidth, but it still requires low-latency communication. When combined with the tree algorithm, efficient parallelization becomes even more difficult.

Distributed-memory parallel computers have been used to run large-scale cosmological simulations, with or without the individual timestep algorithm (Dubinski 1996; Springel et al. 2001). In this case, we use simple spatial decomposition to distribute particles over the processors. This works fine with large-scale cosmological simulations, where the distribution of particles on the large scale is almost uniform. Many structures form from initial density fluctuations, and many small high-density regions develop. Even so, we can still divide the entire system so that the calculation load is reasonably well balanced. In addition, the range of the timesteps is relatively small.

To parallelize the simulation of a single star cluster is much more difficult, because the calculation cost is dominated by a small number of particles in a single, small core (Makino, Hut 1988). Therefore, communication latency becomes the bottleneck, and it is difficult to parallelize the simple direct summation algorithm. As a result, no good parallel implementation of the combination of the tree algorithm and individual timestep algorithm exists. To really accelerate the calculation of a single cluster, we need an approach different from what has been tried.

There are three different approaches to improve the speed of any simulation: a) to use a faster computer, b) to use algorithms with a smaller calculation cost, and c) to improve the efficiency of the algorithm used. Usually, option (a) means to use commercially available fast computers, which at present means distributed-memory parallel computers. An alternative possibility is to develop a computer by ourselves. We

have been pursuing this direction, starting with GRAPE-1 (Ito et al. 1990).

The basic idea of the GRAPE (GRAVity piPE) architecture (Sugimoto et al. 1990) is to develop a fully pipelined processor specialized for calculating the gravitational interaction between particles. In this way, a single force-calculation pipeline integrates more than 30 arithmetic units, which all operate in parallel. In the case of an Hermite time integration, we also need to calculate the first time derivative of the force, resulting in nearly 60 arithmetic operations. This means that we can integrate a large number of arithmetic units into a single hardware with a minimal amount of additional logic.

GRAPE-1 was an experimental hardware with a very short word format (relative force accuracy of 5% or so), and was not really suited for simulations of collisional systems. However, its exceptionally good cost–performance ratio made it useful for simulations of collisionless systems (Okumura et al. 1991; Funato et al. 1992). Also, we developed an algorithm to accelerate the Barnes–Hut tree algorithm using GRAPE hardware (Makino 1991c), and developed GRAPE-1A (Fukushige et al. 1991), which was designed to achieve good performance with the tree code. Thus, the GRAPE approach turned out to be quite effective, not only for collisional simulations, but also for collisionless simulations as well as SPH simulations (Umemura et al. 1993; Steinmetz 1996). GRAPE-1A and its successors, GRAPE-3 (Okumura et al. 1993) and GRAPE-5 (Kawai et al. 2000), have been used by researchers worldwide for many different problems.

In this paper, we discuss GRAPE-6, our newest machine for simulating collisional systems. We briefly summarize the history of the hardwares here.

GRAPE-2 (Ito et al. 1991) adopted the usual 64- and 32-bit floating-point number format, and could be used with Aarseth’s NBODY3 program. After GRAPE-2, we developed GRAPE-3 (Okumura et al. 1993), which was essentially an LSI implementation of GRAPE-1. In GRAPE-1, arithmetic operations were realized by fixed-point ALU chips and ROM chips, and in GRAPE-2 by floating-point ALU chips. Thus, we needed several tens of LSIs to realize a single pipeline. With GRAPE-3, we implemented a single pipeline to a single custom LSI chip, and developed a board with 24 chips. In this way, we achieved a speed of 9 Gflops per board (24 chips each performing 38 operations on 10 MHz clock cycle).

GRAPE-4 (Makino et al. 1997) was similarly a single-LSI implementation of GRAPE-2, or actually that of HARP-1 (Makino et al. 1993), which was designed to calculate force and its time derivative. A single GRAPE-4 chip calculated one interaction in every three clock cycles, performing 19 operations. Its clock frequency was 32 MHz and peak speed of a chip was 608 Mflops.

A major difference between GRAPE-4 and previous machines was its size. GRAPE-4 integrated 1728 pipeline chips, for a peak speed of 1.08 Tflops. The machine was composed of 4 clusters, each with 9 processor boards. A single processor board housed 48 processor chips, all of which shared a single memory unit through another custom chip to handle predictor polynomials. GRAPE-4 chip used two-way virtual multiple pipeline, so that one chip looked like two chips with half the clock speed. Thus, one GRAPE-4 board calculated the

forces on 96 processors in parallel. Different boards calculated the forces from different particles, but to the same 96 particles. Forces calculated in a single cluster were summed up by special hardware within the cluster.

In this paper, we describe the architecture and performance of GRAPE-6, which is the direct successor of GRAPE-4. The main difference between GRAPE-4 and GRAPE-6 is in the performance. The GRAPE-6 chip integrates 6 pipelines operating at 90 MHz, offering a speed of 30.8 Gflop, and the entire GRAPE-6 system with 2048 chips offers a speed of 63.04 Tflop.

The plan of this paper is as follows. In section 2, we describe the overall architecture, and in sections 3 and 4 the details of implementation. In section 5, we discuss the difference between GRAPE-4 and GRAPE-6. In section 6 we discuss the performance. Section 7 is for discussions. Those who are interested in how to use GRAPE-6, but not much in the design details, could skip subsection 2.1, most of section 3 and section 5.

2. The Architecture of GRAPE-6

In this section, we give an overview of the architecture of GRAPE-6. What GRAPE-6 calculates are the following. First, it calculates the gravitational force, its time derivative, and potential, given by equations:

$$\mathbf{a}_i = \sum_j Gm_j \frac{\mathbf{r}_{ij}}{(r_{ij}^2 + \epsilon^2)^{3/2}}, \quad (1)$$

$$\dot{\mathbf{a}}_i = \sum_j Gm_j \left[\frac{\mathbf{v}_{ij}}{(r_{ij}^2 + \epsilon^2)^{3/2}} - \frac{3(\mathbf{v}_{ij} \cdot \mathbf{r}_{ij})\mathbf{r}_{ij}}{(r_{ij}^2 + \epsilon^2)^{5/2}} \right], \quad (2)$$

$$\phi_i = \sum_j Gm_j \frac{1}{(r_{ij}^2 + \epsilon^2)^{1/2}}, \quad (3)$$

where \mathbf{a}_i , $\dot{\mathbf{a}}_i$, and ϕ_i are the gravitational acceleration, its first time derivative, and the potential of particle i ; m_i , \mathbf{x}_i , and \mathbf{v}_i are the mass, position, and velocity of particle i , G is the gravitational constant, and ϵ is the softening parameter. GRAPE-6 hardware assumes $G = 1$. If necessary, the host computer can multiply the result calculated by GRAPE-6 by some constant to use a G other than one. Also note that the potential is calculated without minus sign. The relative position, \mathbf{r}_{ij} , and relative velocity, \mathbf{v}_{ij} , are defined as

$$\mathbf{r}_{ij} = \mathbf{x}_j - \mathbf{x}_i, \quad (4)$$

$$\mathbf{v}_{ij} = \mathbf{v}_j - \mathbf{v}_i. \quad (5)$$

While calculating the force, it also evaluates the distance to the nearest neighbor,

$$r_{\min} = \min_{j \neq i} r_{ij}, \quad (6)$$

and the value of index j , which gives the minimum distance. In addition, it constructs a list of neighboring particles, whose distance squared (with softening, $r_{ij}^2 + \epsilon^2$) is smaller than a pre-specified value h_i^2 .

The position \mathbf{x}_j and velocity \mathbf{v}_j of particles that exert forces are ‘‘predicted’’ by the following predictor polynomial:

$$\mathbf{x}_{j,p} = \frac{\Delta t_j^4}{24} \mathbf{a}_{j,0}^{(2)} + \frac{\Delta t_j^3}{6} \dot{\mathbf{a}}_{j,0} + \frac{\Delta t_j^2}{2} \mathbf{a}_{j,0} + \Delta t_j \mathbf{v}_{j,0} + \mathbf{x}_{j,0}, \quad (7)$$

$$\mathbf{v}_{j,p} = \frac{\Delta t_j^3}{6} \mathbf{a}_{j,0}^{(2)} + \frac{\Delta t_j^2}{2} \dot{\mathbf{a}}_{j,0} + \Delta t_j \mathbf{a}_{j,0} + \mathbf{v}_{j,0}, \quad (8)$$

where $\mathbf{x}_{j,p}$ and $\mathbf{v}_{j,p}$ are the predicted position and velocity; $\mathbf{x}_{j,0}$, $\mathbf{v}_{j,0}$, $\mathbf{a}_{j,0}$, and $\dot{\mathbf{a}}_{j,0}$ are the position, velocity, acceleration, and its time derivative of particle j at time $t_{j,0}$; and Δt_j is the difference between the current time t_j of particle j and system time t , i.e.,

$$\Delta t_j = t - t_j. \quad (9)$$

2.1. Individual Timestep on GRAPE Hardware

Here, we briefly summarize how GRAPE-6 (and GRAPE-4) works with the individual timestep algorithm. For a more detailed discussion, see Makino et al. (1997) or Makino and Taiji (1998).

The time integration proceeds in the following steps:

- a) As the initialization procedure, the host sends all data (position, velocity, acceleration, its first time derivative, mass, and time) of all particles to the GRAPE memory unit.
- b) The host creates a list of particles to be integrated at the present timestep.
- c) For each particle in the list, steps (d)–(g) are repeated.
- d) The host predicts the position and velocity of the particle, and sends them to GRAPE. GRAPE stores them in the registers of the force calculation pipeline. It also sets the current time to a register in the predictor pipeline.
- e) GRAPE calculates the force from all other particles. The positions and velocities of other particles at the current time are calculated in the predictor pipeline.
- f) After the calculation is finished, the host retrieves the result.
- g) The host integrates the orbits of the particles and determines new timesteps.
- h) The present system time is updated and the process returns to step (b).

Here, the key to achieving good performance is to send only particles updated in the current timestep to the GRAPE hardware. Thus, the GRAPE hardware needs to have a memory unit large enough to keep all particles in the system. This is usually not a severe limitation, since even with fast GRAPE hardwares, the number of particles that we can handle with the direct summation algorithm is not very large.

2.2. Top-Level Network Architecture

The top-level architecture of GRAPE-6 is shown in figure 1. It consists of 4 ‘‘clusters’’, each of which comprises 16 GRAPE-6 processor boards (PB), 4 host computers (H), and interconnection networks. These 4 clusters are connected by Gigabit Ethernet. For host computers, we currently use PCs with AMD Athlon XP 1800+ CPU and SiS 745 chipset. Ethernet cards are 1000BT cards with NS 83820 single-chip Ethernet controllers.

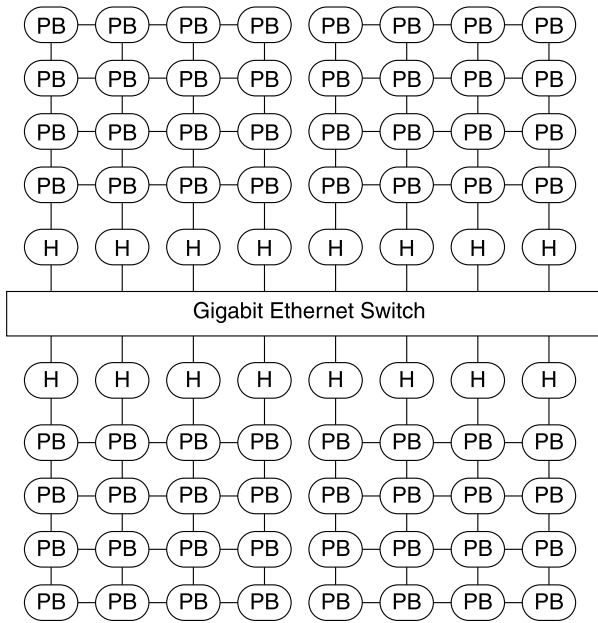


Fig. 1. Top level network structure of GRAPE-6. “H” indicates a host computer and “PB” indicates a processor board.

In the following, we describe how we run parallel programs on GRAPE-6. First, let us concentrate on parallelization within a cluster.

Figure 2 shows one cluster. Four processor boards are connected to a host computer through a network board. Four network boards are connected to each other, so that we can use a cluster as a single unit or as multiple units.

First consider the simplest case, where we just use 4 hosts to run independent calculations. In this case, 4 processor boards connected to a host through one network board calculate the forces on the same set of particles, but from a different set of particles [what we called *j*-parallelism in Makino et al. (1997)]. Each processor board stores different subsets of particles in the particle memory, and calculates the forces on the particles stored in the registers in the processor chips. The partial forces calculated in different boards are sent in parallel to the network board, where they are added together by an adder tree. The host computer receives the summed-up forces. As discussed later, multiple processor chips on one board also have their local memories to store particles. They calculate the forces on the same set of particles, but from different sets of particles. The partial forces are summed up by the adder tree on the processor board. From a logical point of view, there is no difference between a single-board system and multi-board system, as long as we use a single host. We can regard the entire system as just a huge adder tree with processor chips at all leaves.

When all 16 boards and 4 hosts are used as a single unit, the particles are divided to 4 groups, and each group is assigned to one host. Conceptually, the *j*-th board connected to host *i* calculates the force on particles in host *i*, from particles in host *j*. Summation of the partial forces is performed in the same way as in the case of a single-host calculation. The only difference is that the data to be stored

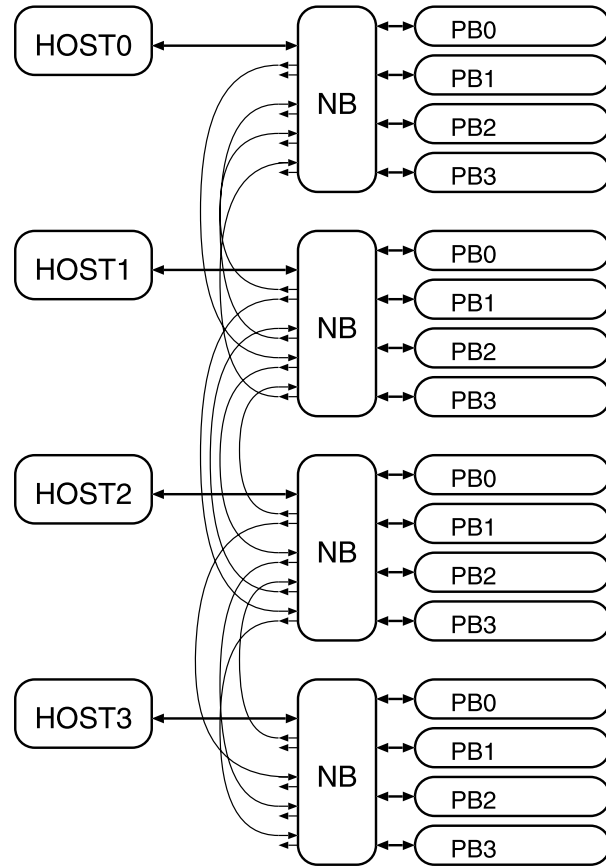


Fig. 2. GRAPE-6 cluster. “NB” indicates a network board.

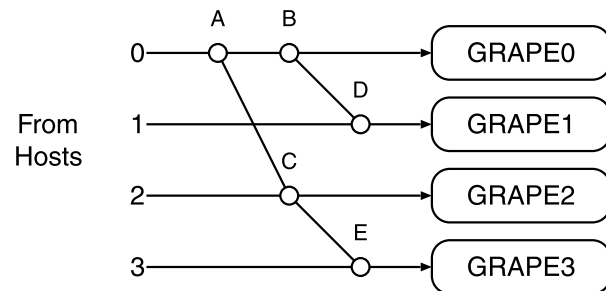


Fig. 3. A 4-input example of a switching network for parallel GRAPE.

in the memory come from other hosts.

In order to allow both single-host and multi-host calculations, the network board must switch between the broadcast mode (for the single-host calculation) and the point-to-point mode (for the multi-host calculation). It would also be useful if we can use two hosts together. In this case, it is necessary to accept two inputs, and to pass each of them to two boards. Thus, we need three operation modes for the network board. One simple way to implement these three modes is shown in figure 3. Here, nodes A and B simply output the inputs from the left-hand side ports to two output ports. Nodes C, D, and E can select one from two inputs. In the case of node C, the selected input is sent to two output ports.

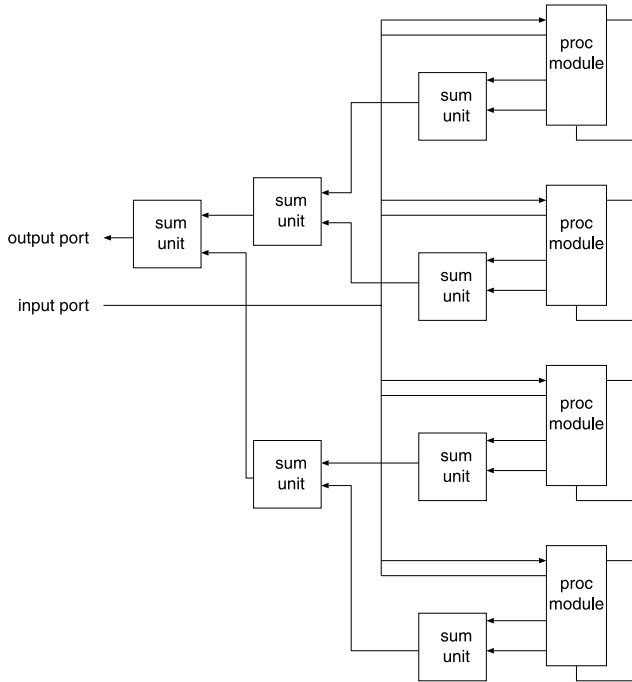


Fig. 4. Structure of the processor board.

This network can be configured in three ways. In the first mode, all nodes select the input from the lower ports in figure 3. In other words, C takes the input from input port 2, D from input port 1, and E from input port 3. In this case, each GRAPE receives data from the input port with the same index. In this mode we can use this 4-GRAPE network as part of a 4-host, 16-GRAPE system. In the second mode, node C selects the input from port 2, while D and E select data from upper input port in figure 3 (nodes B and C for nodes D and E, respectively). In this mode, GRAPEs 0 and 1 receive the same data from port 0. Similarly, GRAPEs 2 and 3 receive data from port 2. In other words, GRAPEs 0 and 1 (and 2 and 3 as well) are effectively bundled together to behave as one system, and we can use this system as a part of the 2-host, 8-GRAPE system. In the third mode, all nodes select upper inputs, thereby sending the data from port 0 to all GRAPEs. In this way, we can use this 4-GRAPE network as a single system connected to one host.

An important character of this network is that its hardware cost is $O(p)$, where p is the number of GRAPE hardwares. Thus, even for very large systems, the cost of the network remains small. By using this hardware network to send data from multiple hosts to processor boards under one host in parallel, we can improve the parallel efficiency quite significantly.

There are many possible algorithms to parallelize the calculation over multiple clusters. Here, we show just one example, which is a generalization of the “copy” algorithm (Makino 2002). In the copy algorithm, each node has the complete copy of the system. At each timestep, each node, however, integrates its own share of particles, which is either statically or dynamically assigned to it. After one step is finished, all nodes broadcast the particles that they updated,

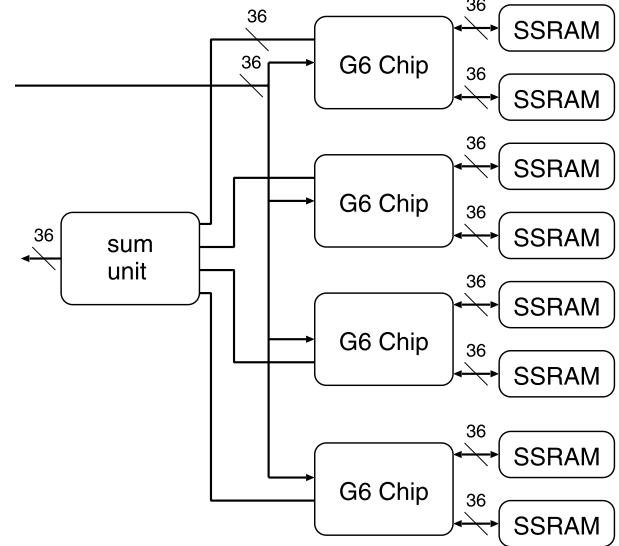


Fig. 5. Structure of the processor module.

so that all nodes have the same updated system. In the case of a multi-cluster calculation, each cluster has a complete copy of the system, which is distributed to 4 hosts. For example, host 0 of cluster 0 and host 0 of cluster 1 have the same data. In the time integration, calculation load is divided between all hosts in the different clusters with the same internal index. After one step is finished, updated data are exchanged again between the hosts in different clusters with the same index. One could use a “ring” algorithm or 2-D algorithm (Makino 2002), but for 4 clusters the difference in the performance is rather small.

In principle, we could extend the network board to form an 8-input, 8-output switch, so that we can use all 64 boards as a “single cluster”. We decided not to do this, since for many scientific applications we will use the system as a correction of single-host systems to run multiple simulations independently. To run multiple calculations, it is more efficient to have larger number of host computers.

2.3. Board-Level Structure

Figure 4 shows the structure of a processor board. It houses 8 processor modules. The processor board has one broadcast network that broadcasts data from the input port to all processor modules, and one reduction network that reduces the results obtained on 32 chips and returns it to the host through the output port.

Each processor module consists of 4 processor chips, each with its memory, and one summation unit. The structure of a processor module is the same as that of the processor board, except that it has 4 processor chips instead of 8 processor modules. Figure 5 shows the structure of a processor module.

The memories attached to one processor chip can store up to 16384 particles. Thus, a single board with 32 chips can handle up to 524288 particles, for a direct summation code with individual timestep. A (4×4) -board cluster can handle up to 2 million particles. If one wants to use more than 2 million particles with direct summation, it is possible to use the ring algorithm (see subsection 5.2). A calculation with 8 million

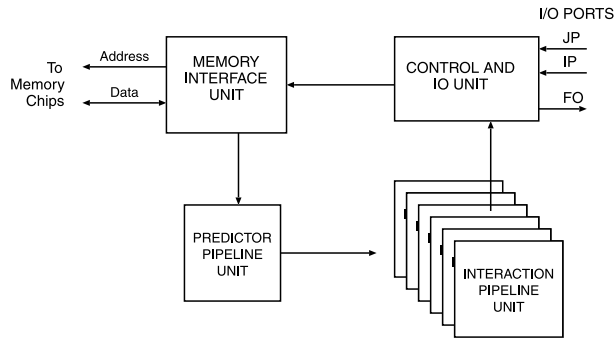


Fig. 6. Block diagram of the processor chip.

particles is theoretically possible on a single cluster with 16 processor boards.

In the next two sections, we present a detailed description of the hardware, in a bottom-up fashion. In section 3, we describe the processor chip and in section 4 the processor board, network board, and interconnection.

3. The Processor Chip

The GRAPE-6 processor chip was fabricated using the Toshiba TC-240 process (nominal design rule of $0.25\ \mu\text{m}$). The physical size of the chip is roughly 10 mm by 10 mm, and is packaged into 480-contact BGA package. It operates at a 90 MHz clock cycle. The power supply voltage is 2.5 V. Heat dissipation is around 12 W at the maximum.

A processor chip consists of six force calculation pipelines, a predictor pipeline, a memory interface, a control unit, and I/O ports. Figure 6 shows an overview of the chip. In the following, we discuss each block in turn.

3.1. Force Calculation Pipeline

The task of the force calculation pipeline is to evaluate equations (1)–(3). It also determines the nearest neighbor particle and its distance. This function is rather convenient for detecting close encounters or physical collisions between particles that require special treatments. For this purpose, the indices of particles that exert forces are supplied to the pipeline.

The indices are also used to avoid self-interaction. The force calculation pipeline has the register for the index of the particle for which the force is calculated, and avoids the accumulation of the result if two indices are the same. This capability is introduced to avoid the need to send particles twice to the memory in the case of the individual timestep algorithm.

With the individual timestep algorithm and the hardwired predictor pipeline, the data of particles which exert forces are evaluated by the predictor pipeline on the chip, while the data for the particle for which the force is calculated is evaluated on the host computer and sent to the register of the force calculation pipeline. These two values are not exactly the same, since the data format and accuracy of the hardware predictor are different from that of the host computer. The GRAPE-4 pipeline does not have logic to use a particle index, and the only way to avoid the self interaction is to make the data exactly the same. To achieve this, for the particles to be updated, we send

Table 1. Arithmetic operations in the force calculation pipeline.

#	Operation	Format	Length	Mantissa
1	$\mathbf{dx} \leftarrow \mathbf{x}_j - \mathbf{x}_i$	fixed	64	...
2	$dr_s^2 \leftarrow \mathbf{dx} ^2 + \epsilon^2$	float	36	24
3	calculation of $r_s^{-\alpha}$, where $\alpha = 1, 3, 5$	float	36	24
4	$\phi_{ij} \leftarrow m_j r_s^{-1}$	float	36	24
5	$\phi_i \leftarrow \phi_i + \phi_{ij}$	fixed	64	...
6	$\mathbf{a}_{ij} \leftarrow m_j r_s^{-3} \mathbf{dx}$	float	36	24
7	$\mathbf{a}_i \leftarrow \mathbf{a}_i + \mathbf{a}_{ij}$	fixed	64	...
8	$d\mathbf{v} \leftarrow \mathbf{v}_j - \mathbf{v}_i$	float	36	24
9	$s \leftarrow d\mathbf{v} \cdot \mathbf{dx}$	float	32	20
10	$j_1 \leftarrow \mathbf{dx} \cdot 3sm_j r_s^{-5}$	float	32	20
11	$j_2 \leftarrow d\mathbf{v} \cdot m_j r_s^{-5}$	float	32	20
12	$\dot{\mathbf{a}}_{ij} \leftarrow j_1 + j_2$	float	32	20
13	$\dot{\mathbf{a}}_i \leftarrow \dot{\mathbf{a}}_i + \dot{\mathbf{a}}_{ij}$	fixed	32	...

the predicted data at the current time to the memory as well as the registers. This means that we have to send j -particles twice per timestep. With the index-based approach, we need to send j -particles only once per timestep, resulting in a significant reduction in the total amount of communication.

For GRAPE-6 pipeline, we have adopted the 8-way VMP (virtual multiple pipeline, Makino et al. 1997), in which a single physical pipeline serves as eight virtual pipelines, calculating the forces on 8 different particles. In this way, we can reduce the requirement for memory bandwidth by a factor of 8, since all VMPs (and also physical multiple pipelines on a chip) calculate the forces from the same particle.

In the physical implementation of the pipeline, we have adopted several different number representations, depending on the required accuracy. For input position data, we use the 64-bit fixed-point format. The reason that we used the fixed-point format here is to simplify the hardware. An additional advantage of using the fixed-point format is that the implementation of the periodic boundary condition is simpler than that in the case of the floating-point data format (Fukushige et al. 1996).

After first subtraction between two position vectors, the result is converted to a floating-point format with a 24-bit mantissa. Here, the floating-point format is preferred, since otherwise we need very large multipliers.

For the final accumulation, we return to the 64-bit fixed-point format, again to simplify the hardware. Here, we specify the scaling factor for each particle, so that we can calculate forces with very different magnitudes, without causing an overflow or underflow.

The pipeline for the calculation of the time derivative is designed in a similar way, but with a 20-bit mantissa for intermediate data and the 32-bit fixed-point format for the final accumulation. Since the time derivative is one order higher than the force, the required accuracy is lower.

Figure 7 shows a block diagram of the pipeline. It consists of arithmetic units to perform the operations listed in table 1. We briefly discuss each operations below.

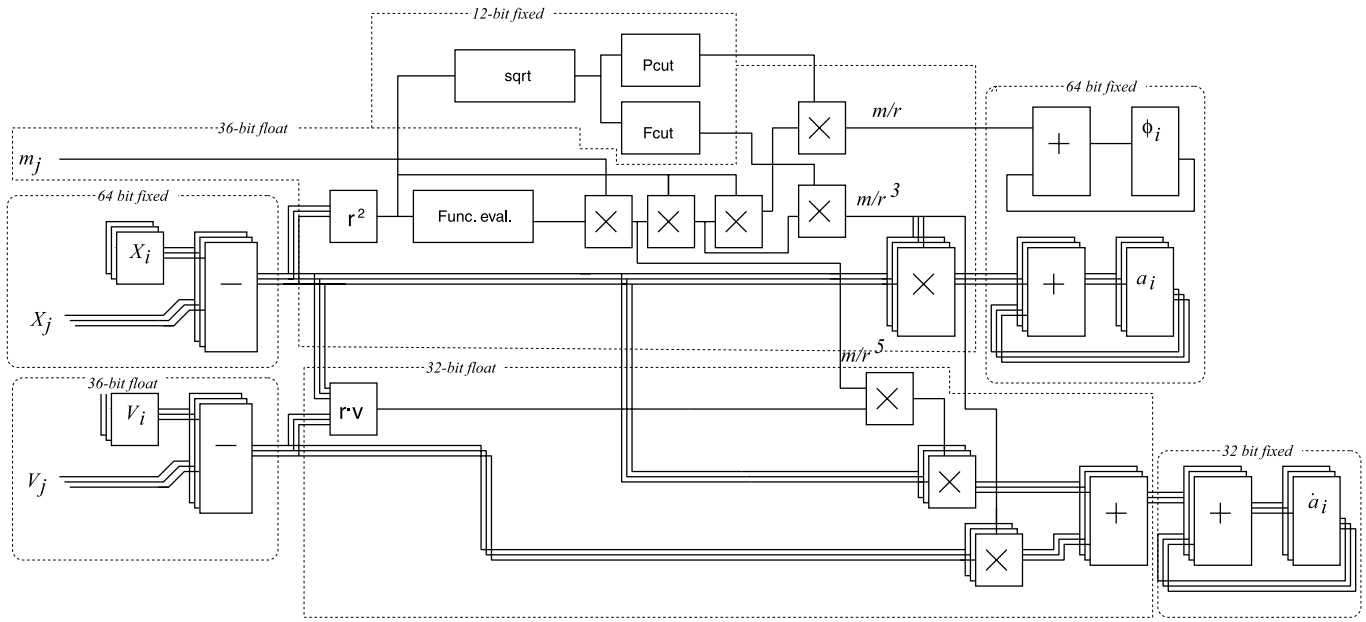


Fig. 7. Block diagram of the force calculation pipeline.

3.1.1. $d\mathbf{x} \leftarrow \mathbf{x}_j - \mathbf{x}_i$

The position data are expressed in the 64-bit 2's complement fixed-point format. The result of subtraction is then converted to the floating-point format. The floating-point format used here consists of the sign bit, the zero bit, 10-bit exponent and 24-bit mantissa. The sign bit expresses the sign (one denotes a negative number). The zero bit indicates if the number is zero or not (one indicates that the expressed value is zero). In the standard IEEE floating-point format, a zero value is expressed in a special format (both exponent and mantissa are zero). This convention is useful to achieve the maximum accuracy for a given word length. However, using zero bit is more cost-effective in the internal expression for the hardware, since the logic to handle a zero value is greatly simplified.

For the result of subtraction, the range of the exponent is 6 bits. We use a biased format, and extend the exponent to 10 bits. For all floating-point operations, we use 10-bit exponents, to avoid overflows in intermediate results (in particular for r^{-5}).

The length of the mantissa is 24 bits, with the usual "hidden bit" at MSB (most significant bit). For the rounding mode, we use the "force-1" rounding, with a correction to achieve unbiased rounding. With the "force-1" rounding, we always set the LSB of the calculated result (after proper shifting) to be one, regardless of the contents of the field below LSB. Thus, if the LSB is already one, the result is rounded toward zero, and if the LSB is zero, the result is rounded toward infinity. Thus, this rounding gives an almost unbiased result.

However, in this simple form this rounding is still biased, since the treatment for the case where all bits below LSB are zero is not symmetric. Consider the following example, where we use 4 bits for the mantissa, and the calculated result is in 8 bits (for example with multiplication). If the result is 10010000 in the binary format, it is rounded to 1001. If the

result is 10000000, it is also rounded to 1001. Thus, out of 32 possible combinations of LSB bit and 4 bits under LSB, for 16 cases the rounding is upward, 15 cases downward, and one case no change. This gives a slight upward bias for the rounded result.

One way to remove this bias is not to force one if all bits below LSB are zero. This can be implemented with a rather simple logic, and we use this method with all floating-point arithmetic units used in GRAPE-6.

Compared to the usual "round-to-the-nearest-even" rounding, this bias-corrected rounding is significantly easier to design and test. In particular, there is no need for the conditional incrementer that would be necessary with the usual nearest rounding. Of course, this simpler design does not mean a smaller number of gates, since our rounding requires the length of the mantissa to be longer than that for the nearest rounding by 1 bit. However, it is also true that the additional number of gates is rather small, because we do not need the conditional incrementer.

3.1.2. $dr_s^2 \leftarrow |d\mathbf{x}|^2 + \epsilon^2$

These are realized with the usual floating-point multipliers and adders. The design of the adder used here is simpler than that of the general-purpose floating-point adders, since we know that both operands are positive. This means that the result of the addition of the two mantissas is always greater than the larger of the two operands, and we need to shift the result at the maximum by one bit. With the general-purpose adder, if two operands have similar magnitudes and different signs, the result of the addition can be much smaller, and we need a shifter with the capability to shift up the result by up to the length of the mantissa, itself.

3.1.3. Calculation of $r_s^{-\alpha}$

Here, we follow the design of the GRAPE-4 chip, where we use a segmented second-order polynomial to calculate $r_s = r_s^{-5}$. We then multiply r_s by m_j , and then by r_s^2 twice to obtain $m_j r_s^{-3}$ and m_j/r_s .

To calculate r_s^{-5} , we first normalize r_s^2 to the range of $[1/4, 1)$. In other words, r_s is expressed as $2^{2a+b} \cdot c$, where a is an integer number, b is either 0 or 1, and c is in the range of $[1/2, 1)$. The “exponent” a is multiplied by -5 to obtain the exponent of the resulting r_s^{-5} .

We use a table with 512 entries to obtain the coefficients for the polynomial. This table accepts b and eight MSB bits (excluding the hidden bit) as the input address. The output of the table consists of the coefficients for the second-order term (12 bits), first-order term (18 bits), zeroth-order term (24 bits), and exponent (3 bits). Note that the calculated result is always smaller than the zeroth-order term, since both the first and second derivatives have minus signs. Therefore, the MSB of the calculated result can turn to zero, even though MSB of the zeroth-order term is always one. In this case, we need to shift the result by one bit, and adjust the exponent by one. This adjusted exponent is then added to a previously calculated exponent to obtain the exponent of the final result.

3.1.4. $\phi_{ij} = m_j r_s^{-1}$

As described in the previous subsection, we actually calculate $m_j r_s^{-5}$ and then multiply it by r_s^2 twice. These are usual floating-point multiplications, with bias-corrected force-1 rounding.

3.1.5. $\phi_i = \phi_i + \phi_{ij}$

The potential is accumulated in the 64-bit fixed-point format. The pairwise potential, ϕ_{ij} , which is obtained in the floating-point format, is shifted before addition according to the shift length, $e_\phi - s_{\phi,i}$, where e_ϕ is the value of exponent of the pairwise potential, ϕ_{ij} , and $s_{\phi,i}$ is the scaling coefficient for the potential of particle i . Note that the coefficient, $s_{\phi,i}$, is specified on a per-particle basis, and we can specify different values for 48 virtual pipelines. This coefficient should be calculated from a reasonably good estimate of the total potential of particle i , to avoid both overflow and underflow during calculation.

3.1.6. $\mathbf{a}_{ij} = m_j r_s^{-3} \mathbf{dx}$

These are usual floating-point multiplications.

3.1.7. $\mathbf{a}_i = \mathbf{a}_i + \mathbf{a}_{ij}$

Here, we use the same design as that for accumulation of the potential. The scaling coefficient is common for all three components.

3.1.8. $d\mathbf{v} \leftarrow \mathbf{v}_j - \mathbf{v}_i$

This and the remaining operations to be discussed in this section are all for the time derivative of the forces. For these operations, we use a number format with a 20-bit mantissa. For this first subtraction, the mantissa of the input is 24 bits, and the result is given with the 20-bit mantissa.

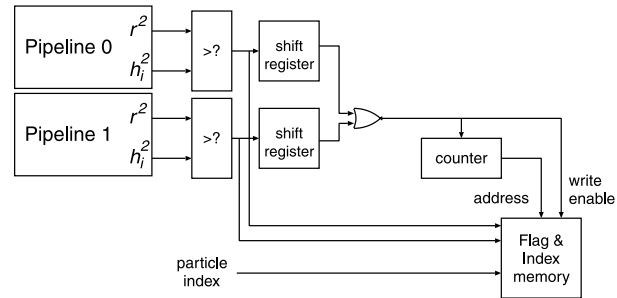


Fig. 8. Block diagram of the neighbor list unit.

3.1.9. $s = d\mathbf{v} \cdot \mathbf{dx}$

This is an inner-product of two vectors in three dimensions. The mantissa of \mathbf{dx} is first truncated to 20 bits.

3.1.10. $\hat{\mathbf{a}}_i = \hat{\mathbf{a}}_i + \hat{\mathbf{a}}_{ij}$

As in the case of the potential and the force, we use a fixed-point format with a scaling coefficient for the final accumulation. Instead of the 64-bit format, however, here we use a 32-bit format, since only the contributions from nearby particles are important for the time derivative.

3.1.11. Cutoff Unit

In figure 7, there are three boxes in the “12-bit fixed” format region. These are used to implement the Gaussian cutoff of the $1/r$ potential, to be used with the Ewald summation method to calculate the gravitational force with a periodic boundary condition. The details of the operation of these boxes will be described elsewhere, along with a discussion of the performance and accuracy of the Ewald method on GRAPE-6. In this paper, we can regard these two boxes, “Pcut” and “Fcut”, as just boxes with constant (unit) outputs.

3.2. Neighbor List Unit

The neighbor list unit of GRAPE-6 chip is essentially the same as that of the GRAPE-4 board. It consists of two memory units, one for the indices of the j -particles and the other for flags to indicate (virtual) pipelines. One neighbor list unit serves 16 virtual pipelines (two physical pipelines). Thus we integrated three units to one chip. One neighbor list unit can store up to 256 neighbor particles.

Figure 8 shows one neighbor list unit. Each pipeline has registers (for each of the virtual pipelines) for the neighbor radius squared h^2 , and if the distance to the current j -particle is not larger than the neighbor radius, a flag is asserted. This flag is stored to a shift register. Once per every eight clock cycles, this shift register contains eight flags from different VMPs for the same j particles. At this cycle, if any of the 16 flags from 16 virtual pipelines is asserted, the index of the current j -particle and the flags, themselves, are written to the memory.

3.3. Predictor Pipeline

The predictor pipeline evaluates the predictor polynomials expressed in equations (7) and (8). As stated earlier, we use 8-way VMP for the force calculation pipeline. Therefore, the predictor pipeline can use eight clock cycles to produce the

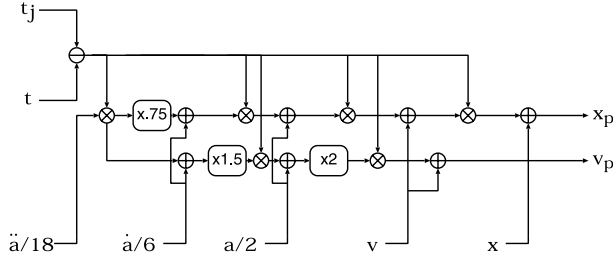


Fig. 9. Block diagram of the predictor unit.

predicted position and velocity of one particle. To take advantage of this fact, we implement one pipeline, which processes the x , y , and z components sequentially. In principle, we could further reduce the hardware size by using one pipeline for both the position and the velocity. We do not adopt this approach, since the circuit size for the predictor pipeline is already a small fraction of the total size of the chip.

In the design of the predictor pipeline, we have tried to minimize the amount of data to express the predictor for one particle, since it directly affects the time needed for communication and the number of wires needed between the memory chips and the processor chip.

With GRAPE-4, the predictor data for one particle is expressed in 19 32-bit words (2 for time, 6 for position, 3 for each of velocity, acceleration, and time derivative of acceleration, 1 for mass, and 1 for memory address). With a similar format, GRAPE-6 predictor would need 23 words, since we need one more word for the particle index and three more for the second time derivative of the acceleration. In many applications, inclusion of the second derivative improves the accuracy rather significantly.

To reduce the data length, we have adopted the following two methods. First, for the particle time, instead of sending the time itself, we send the location of the bit below which the current system time t can be different from the particle time t_j . In this way, we can reduce the number of bits to express time from 64 to just 7.

Second, we use a block floating-point format with the mantissa length optimized for each of the predictor coefficients. Thus, we use 32, 20, 16, and 10 bits, for the velocity, acceleration, and first and second time derivatives, respectively.

With these two changes, we can make one predictor data to be expressed in 16 32-bit words. Thus, we can use a 64-bit memory bus with the clock speed being the same as that for the pipeline, to supply one particle data in every 8 clock periods.

Figure 9 shows a block diagram of the predictor unit.

The predictor pipeline performs the following operations:

- a) $\Delta t \leftarrow t - t_j$,
- b) $p_1 \leftarrow \Delta t \cdot (a^{(2)}/18)$,
- c) $p_2 \leftarrow p_1 \cdot 0.75$,
- d) $p_3 \leftarrow p_2 + (\dot{a}/6)$,
- e) $p_4 \leftarrow p_3 \cdot \Delta t$,
- f) $p_5 \leftarrow p_4 + (a/2)$,
- g) $p_6 \leftarrow p_5 \cdot \Delta t$,
- h) $p_7 \leftarrow p_6 + v$,
- i) $p_8 \leftarrow p_7 \cdot \Delta t$,

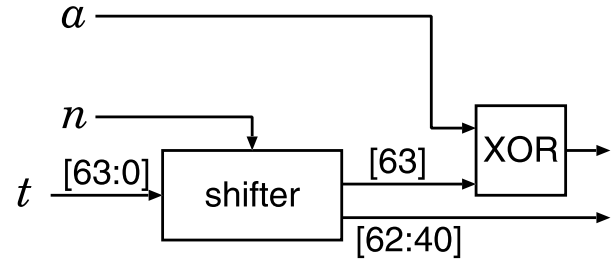


Fig. 10. Block diagram of the logic to handle subtraction of the time.

- j) $x_p \leftarrow x + p_8$,
- k) $q_1 \leftarrow p_1 + (\dot{a}/6)$,
- l) $q_2 \leftarrow q_1 \cdot 1.5$,
- m) $q_3 \leftarrow q_2 \cdot \Delta t$,
- n) $q_4 \leftarrow q_3 + (a/2)$,
- o) $q_5 \leftarrow q_4 \cdot 2$,
- p) $q_6 \leftarrow q_5 \cdot \Delta t$,
- q) $v_p \leftarrow v + q_6$.

Here, p_i is the output of i -th arithmetic unit of the predictor pipeline for the position, and q_i is that for the velocity.

In the following, we describe operations a, b, c, d, j, and q. Other operations are simple fixed-point addition, multiplication, or multiplication by a constant implemented in a way similar to that for operations b, d, and c.

3.3.1. $\Delta t \leftarrow t - t_j$

The current system time t is expressed in the 64-bit fixed-point format. The particle time t_j is expressed by the bit location n , above which t and t_j are the same. This location is the location of MSB of the timestep Δt_j . Consider the following example. If $t_j = 0.5$ and $\Delta t_j = 0.125$, the current time t must be in the range of $[t_j, t_j + \Delta t_j]$, i.e., $[0.5, 0.625]$. In this case, $t - t_j$ can be calculated by simply masking all bits equal to or higher than MSB of Δt_j (i.e., 0.125 and above). This works for any value of t in the range $[0.5, 0.625]$. However, if $t = 0.625$, this procedure returns 0, but the correct value is 0.125. This is simply because we masked the bit which represents the exact value of Δt_j . This problem can be solved by supplying the value of t_j at that bit. Unless t is equal to $t_j + \Delta t_j$, the values of this bit for t and t_j are the same. In this case, the corresponding bit of the resulting Δt must be 0. However, if t is equal to $t_j + \Delta t_j$, the values of this bit for t and t_j are different. In this case, MSB of the result must be one. Thus, by taking XOR of the two input bits, we can determine the MSB value of the result. Figure 10 shows the actual circuit. Here, a is the value of the bit of t_j which corresponds to the non-zero bit of Δt_j and n is the location of that bit. The result is expressed in a 24-bit unsigned fixed-point format. Here, the rounding is simple rounding to zero. This can cause a very small bias in the predicted position, if the timestep of the current blockstep is very small and the timestep of the predicted particle is large. In this case, however, the error in the prediction does not degrade the accuracy. Therefore we do not perform a rounding correction here.

3.3.2. $p_1 \leftarrow \Delta t \cdot (a^{(2)}/18)$

Both inputs are supplied in a 10-bit fixed-point format. Here, we use the sign–magnitude format, instead of the usual 2’s complement format, to simplify the design of the multiplier. Note that Δt is supplied in the 24-bit format. Therefore, we need to truncate it to the 10-bit format, using the bias-corrected force-1 rounding discussed earlier. The result is also rounded to the 10-bit format.

3.3.3. $p_2 \leftarrow p_1 \cdot 0.75$

This multiplication by a constant is achieved by adding $p_1/2$ and $p_1/4$. These two values can be calculated by shifting them to the right by one bit and two bits, respectively. These shiftings, in hardware, require just wiring and no logic. Thus, this multiplication is actually implemented by a single adder. Here, we do not round the result, since the effect of the error in the multiplication in the predictor is usually very small.

3.3.4. $p_3 \leftarrow p_2 + (\dot{a}/6)$

Here, p_2 is in the 10-bit sign-and-magnitude format, and $(\dot{a}/6)$ is in a 16-bit format. Thus, we first extend p_2 to 16 bits. If two inputs have different signs, we need to determine which one is larger. We do this by calculating both $a - b$ and $b - a$. If $a - b$ does not cause an overflow, we can see that $a \geq b$, and use $a - b$ as the result. The sign of the result is the same as that of a . This circuit is rather complicated and larger than that for the 2’s complement format. However, the gate count is practically negligible.

3.3.5. Operations (e)–(i)

All of these are the usual fixed-point addition or multiplication, implemented in the same way as operations (b) and (d).

3.3.6. $x_p \leftarrow x + p_8$

Here, we add two numbers in different formats. One is x in the 64-bit 2’s complement fixed-point format. The other is p_8 in the floating-point format with a sign, exponent and mantissa. Since we do not perform any normalization during the calculation of p_8 , the mantissa is not normalized. This means that we do not use the hidden bit for p_8 . We first shift p_8 according to the value of the exponent of the velocity, and then add it to (or subtract it from) x according to the sign bit.

3.3.7. Operations (k)–(p)

These operations are implemented in the same way as similar operations for the position predictor pipeline are implemented.

3.3.8. $v_p \leftarrow v + q_6$

This is essentially the same addition as used in other operations, but here we post-normalize the result. For the output format we use a mantissa with the hidden bit.

3.4. Memory Interface

The memory interface has two functions. The first one is to write the data sent from the host, and the second one is to read the memory during a calculation.

The data of one particle is packed into 16 32-bit words. A data packet sent from the host consists of two control

Table 2. GRAPE-6 chip input port signal definition.

Signal	Width	Description
DATA	36	32 bit data with 4 bit parity
WE	1	write enable

words and this 16-word data. The first control word contains following three fields: command code (2 bits), chip identity (10 bits), and chip identity mask (10 bits). The second word is the starting address in the memory for the particle data.

The chip identity field is used to select the chip that actually stores the particle data. With the design of GRAPE-6, all chips on one board, or on multiple boards connected to the same host, receive the same data from the host. We, however, have to let different chips calculate the forces from different particles, and this can be achieved by specifying, in the particle data packet, the identity of the chip that actually store the data. When a chip receives one j -particle data packet, it writes the data to the memory only if the chip identity field of the packet (masked by the identity mask) is the same as its identity register (also masked by the identity mask). The identity register itself must be all different on different chips. How we achieve this is discussed in the next subsection. The identity mask field is usually all ones.

The memory interface is designed to control two SSRAM (synchronous static random-access memory) chips with a 36-bit data width. All of the signal lines drive only one chip, so that we can minimize the signal length. Using a combined data width of 72 bits, we implemented ECC (SECDED or single error correct and double error detect) for data received from the memory.

The memory interface is programmable, in the sense that practically all access latencies can be adjusted by writing to on-chip registers. Thus, we can use almost any type of SSRAM with different access timings.

During the calculation, both memory chips output data at every clock cycle. The memory address counter is initialized to $8N$, where N is the number of particles, and decremented at each clock. For writing data, we use a slower access, where we write two SSRAM chips at alternate clock cycles. In this way, we can reduce the switching noise and can also relax the timing requirement for the data bus.

3.5. I/O Ports and Handshake Protocol

Tables 2 and 3 show the signal definition for the input and output ports. Both ports operate on the clock with a frequency 1/4 of that of the internal logic and memory interface. As a result, the communication bandwidth is rather limited. However, the electrical design of the board is easier with a lower clock speed. Also, with the 32-bit data width, we can still achieve a data-transfer speed of around 100 MB s^{-1} , which is fast enough to match with the speed of the PCI bus of the host computer.

The input port is very simple, with data lines and a single write-enable line. The chip actually has two input ports, one dedicated to the data sent to the memory (we call this the JP port), and the other for everything else (the IP port). On the

Table 3. FO port signal description.

Signal	Direction (I/O)	Description
D0-D35	O	data (4 bits for parity)
VD	O	valid data
ND	O	new data
STS	O	status
ACTIVE	O	if 0, chip is unused
WD	I	wait data

JP port, the data of one particle consists of 18 32-bit words, and the control logic handles this 18-word packet. The IP port is a general-purpose port. It accepts variable-length data packets. The first word of the packet is the starting address of the on-chip register. The second word is the number of data words to follow, and the remaining words are all data.

The output port is more complicated, because we need to implement flow control. The reason we need flow control is that for some data, for example for the neighbor list, the host must receive the data directly from all chips. In the case of the force, all chips output the results synchronously and the onboard reduction network reduces the data on the fly. However, the neighbor list data has to be transferred to the host without any reduction.

It is possible to read the data of neighbor particles from each chip without using hardware flow control by letting the host computer send the commands to each chip sequentially until it receives all data. In this case, the processor chip itself does not need any flow control. However, this procedure would be rather slow, since the host has to set up the DMA transfer many times. Therefore, we chose to let the host send the command to all chips. The reduction network takes care of the flow control. In table 3, the WD signal is used for flow control.

When the WD signal is asserted, the chip stops sending new data. When the chip sends new data, it asserts both VD and ND signals. The VD signal is asserted as long as the data is valid, but ND is asserted only when the data is actually updated. The STS line is a special signal which tells whether the force calculation pipeline is working or not. The ACTIVE signal is used to indicate defective chips. The output of this pin is programmable from the host, and if ACTIVE is negated, the reduction network ignores the output from the chip.

4. Processor Board and Network Hardware

4.1. Processor Module and Processor Board

Figures 4 and 11 show the processor board. A single board houses 32 processor chips. Logically, the design of the board is rather simple. The input data is broadcasted to all chips, and the output data of the chips are reduced through a reduction network.

The nodes of the reduction network are made of FPGA chips. It has two operation modes, a reduction mode and a pass-through mode. In the reduction mode, it receives data from lower-level nodes (either the processor chips or lower-level FPGA nodes), and performs reduction. Since one particle data consist of force, potential, time derivative of the force, the

**Fig. 11.** Processor board.

distance and index of the nearest neighbor particle, and status flags, the operation of the reduction ALU needs to change according to the data type, and is controlled by a sequencer.

In the pass-through mode, a node sends data received from the lower level node without applying any operation. Since multiple lower-level nodes might try to send data simultaneously, every node controls the WD signal (which is also implemented in a node FPGA as well) so that only one chip (or node) actually sends data at one time. When one chip (or node) indicates the end of the data by negating the VD signal, the node negates the WD signal for the next chip to start receiving data from that chip.

As can be seen in figure 11, one processor board is designed to house up to eight processor “modules”. A single module houses 4 processor chips, 8 SSRAM chips, and an FPGA chip which realizes a 4-input, 1-output reduction tree. We made this division between the board and module, to make the manufacturing easier. With this separation, all BGA chips (with large number of pins) are mounted on small-size module boards. Thus, the rate of the soldering error should be lower, compared to the case where we mount them on large boards. In addition, if there is an error, only a module with 4 chips would be defective. Of course, having to connect the board and module through a connector increases the probability of a failure, but we expected that the failure rate of the connector would be significantly lower than the failure rate of the soldering (which turned out to be the case)

The tree nodes are implemented using Altera ACEX series chips. In the lowest level (processor module level), we used EP1K50A chips in 484-connect BGA packages. This chip implements a four-input node. Higher levels are implemented on EP1K30A chips in 208-pin QFP packages. This chip implements a 2-input node. These nodes are on the processor board.

The processor board is an 8-layer standard PCB. The processor module board is an 11-layer board with inner via holes. The FPGA and processor chips are mounted on the top side, and SSRAM chips on the bottom side. By this layout, we can minimize the wire length between SSRAM chips and processor chips, and still achieve a rather high packaging density. We use 4 Mbit SSRAM chips. Two SSRAM chips

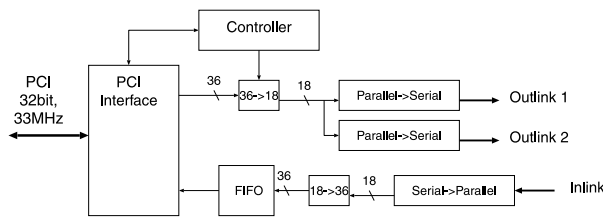


Fig. 12. Host interface board.

connected to a processor chip can store up to 16384 particles. One board can store up to 524288 particles.

The SSRAM chip we chose requires 2.5 V power supply for I/O and 3.3 V for the core. Both the processor board and the module board have separate power planes for both 2.5 and 3.3 V power supplies.

Though the chip has separate ports for j -particles and other data, for the board we decided to use a common data line to simplify the design and reduce the manufacturing cost.

Currently, the core of the processor chip operates on a 90 MHz clock, and the I/O part on a 22.5 MHz clock. The reduction network and other logics of the control board also operate on a 22.5 MHz clock.

For the board-board connection, we use a semi-serial LVDS signal. We use 4-wire (3 for signals and 1 for transmission clock) chipset, which performs 7 : 1 parallel-serial conversion. Since our basic transfer unit is a 32-bit word, we use two cycles of this chipset to transmit one data. Thus, the chipset operates on a 45 MHz clock, and the signal lines operate at a data rate of 315 MHz. For conversion between the 22.5 MHz data rate of the board logic and the 45 MHz data rate of the LVDS chipset, we use additional FPGA chip.

With this LVDS chipset, the receiver chipset itself is driven by the clock signal that comes with the data. In order to allow the two boards connected to a link to operate on independent clocks, we add FIFO chips after the data rate is reduced to 22.5 MHz.

The physical form factor of the card is that of an 8U Eurocard (with the length of 400 mm). For the backplane connection, we use connectors designed for Compact PCI cards. The power supply is also from a backplane bus, through special power connectors.

It is possible to connect a single processor board directly to the host through the host interface card, without using the network card. For this purpose, the processor board also has connectors for twisted-pair cables for the LVDS interface. These connectors are standard RJ-45 modular jacks widely used for 10/100/1000BT Ethernet connection. Standard category 5 (or enhanced 5) cables can be used for connections.

For LVDS interface chips, we use SN75LV85 and SN75LV86A chips from Texas Instruments.

4.2. Host Interface Card

Figure 12 shows a block diagram of the host interface card. It is a standard (32-bit, 33 MHz) PCI card. To transfer data from the host to GRAPE-6, the host sets up the data to be transferred in its memory and lets the PCI interface chip on the interface card perform DMA transfer. The data received

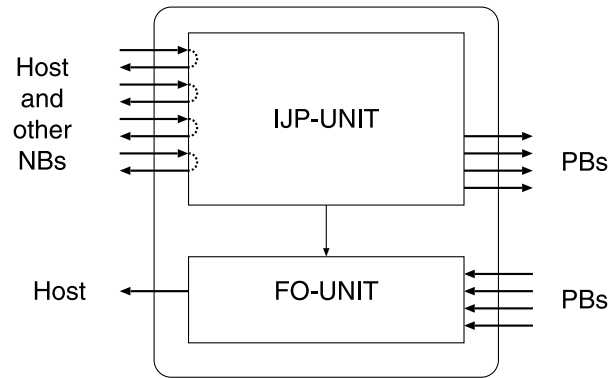


Fig. 13. Network board.

by this DMA transfer is sent directly through the output link. In the design of the host interface card we implemented two output ports so that they can separately supply data to the JP and IP ports. As stated earlier, we decided to use only one port for the processor board. Therefore, the second output port of the interface card is not used.

The input port is more complicated, with an FIFO memory to store the received data. This FIFO memory is necessary, since we cannot guarantee the response time of the host operating system to a DMA request from the interface card. We need to have the memory large enough to avoid any possible overflow.

For the PCI interface, we use the 9080 chip from PLX technology.

4.3. Network Board

Figure 13 shows a block diagram of the network board. It has two basic functions. One is to broadcast (or multicast) data received from the host (or possibly higher-level network boards) to the processor boards (or lower-level network boards). This part is shown as IJP-UNIT. The other is a reduction network for the calculated result, shown as FO-UNIT. The reduction network is exactly the same as that on the processor board, except for the fact that the interface to the module board is replaced by the interface for the processor board (with LVDS link chipset).

The IJP-UNIT has 4 input ports. One of them is a special port designed to connect to the host. The other three ports are designed to accept data from other network boards (see figure 2). Each of the four input ports has a “copy” output, shown on the left-hand side of the unit, so that we can cascade multiple network boards.

Figure 14 shows a block diagram of the multicast network. The boxes in the center of the figure are all data buffers with an output enable control input, which realize the multicast network. Note that this structure implements a network logically equivalent to what is shown in figure 3.

The control input for these buffers is supplied from the control logic implemented on the FPGA for a 45 MHz–22.5 MHz data rate change. This FPGA integrates a sequencer to decode IP/JP port data packets, which reacts to the address space assigned to the network board.

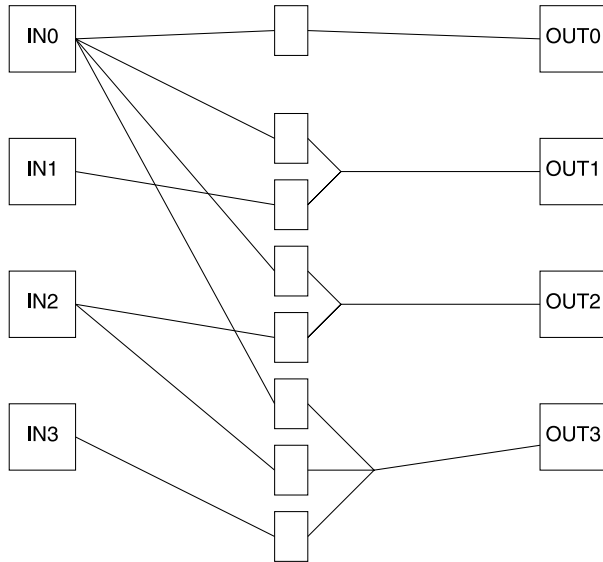


Fig. 14. Physical implementation of the multicast network.

4.4. Packaging and Power Distribution

In the standard configuration, eight processor boards and two network boards are installed in a card rack with a special backplane for the LVDS link. The network board is a single-height unit, but the processor board occupies a two-unit height, to allow sufficient airflow.

To both of the network boards and processor boards, electrical power is supplied through backplane connectors. However, in our present packaging, each processor board has its own power supply unit. A power supply unit accepts DC 330 V input, and supplies DC 2.5 V and 3.3 V. The DC 330 V power is generated by another power unit from a three-phase AC power line. For all of these power units, we use products from Vicor.

We chose Vicor product primarily to reduce the response time of the power supply to the change in power consumption by the boards. One advantage of the CMOS logic is that it consumes power only when the logic state changes. This means that even though we pay absolutely no effort to reduce the power consumption of the chip, its power consumption almost halves when the pipeline is not active.

This “feature” of the chip is rather good from the point of view of the running cost of the machine, but causes a rather serious problem to the power supply. The typical response timescale of a switching power supply unit is on the order of one millisecond. On the other hand, GRAPE-6 switches between calculation and idle (or communication) states in about one millisecond. This means that the response time of the power supply is too long to compensate for a change in the load between the calculation state and the idle state, and the supply voltage becomes rather unstable. Thus, we had to look for power supplies with a relatively short response time. For switching power supplies, a short response time means a high operating frequency, and Vicor products had the highest frequency among commercially available power units.

Even with high-frequency power supplies, the response time

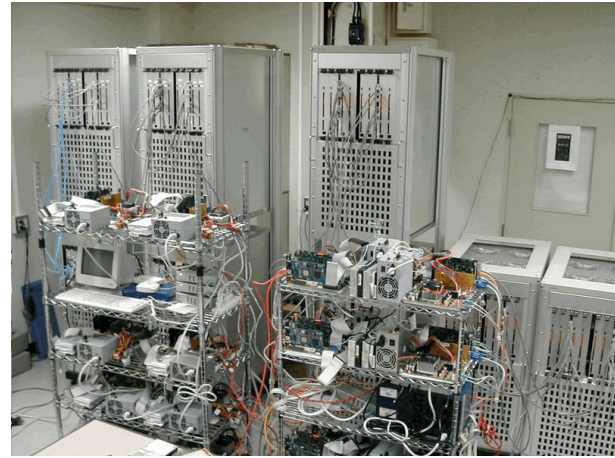


Fig. 15. The 64-board, 4-cluster GRAPE-6 with the racks for the host computers in front.

was still on the order of 100 microseconds, and the only way to stabilize the power supply was to add large bypass capacitors. We attached capacitors with a total capacitance of about 0.1 F to the 2.5 V power line of each processor board. We could not use the usual aluminium electrolytic capacitors because their internal resistance (equivalent series resistance, ESR) is too large. We used low-ESR electrolytic capacitors from Sanyo to meet our need.

In hindsight, it would probably have been better to design a small switching power supply unit integrated into the processor module, since such a power supply unit, which is used on every motherboard for PCs, is inexpensive and highly reliable.

Figure 15 shows the complete GRAPE-6 system consisting of five racks (three with two subracks and two with one subracks), with 16 host computers in front of them. Host computers are Linux-running PCs, with AMD Athlon XP 1800+ processors and ECS K7S6A motherboards. They are connected with Gigabit Ethernets. The total power consumption of the system is around 40 kW, when in full operation.

5. Differences between GRAPE-4 and GRAPE-6

As described in the previous sections, the architecture of GRAPE-6 is quite different from that of GRAPE-4, even though it is the direct successor of GRAPE-4 for essentially the same goal. In this section, we describe what design changes were made, and why.

5.1. Differences in the Semiconductor Technology

The primary difference is that for the GRAPE-6 processor chip we use the 0.25 μm design rule, while with GRAPE-4 we used the 1 μm design rule. This difference with an additional advance in wiring enables us to integrate a roughly 20-times larger number of transistors, with a 3–4 times faster clock speed. Thus, roughly speaking, a single GRAPE-6 chip offers a speed two orders of magnitude higher than that of GRAPE-4.

This large advance, however, implies that almost every design decision had to be changed. In the following, we summarize the changes made.

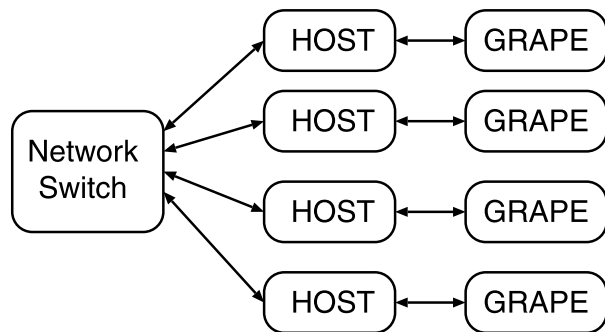


Fig. 16. A simple parallel-host, parallel-GRAPE system.

5.2. Host Computer and Overall Architecture

In GRAPE-4, 4 clusters were connected to a single host, sharing one I/O bus. For the peak speed of 1 Tflops, a single host was still okay for simulations with a large number of particles (10^5 and larger), and communication through a single I/O bus was also okay.

With GRAPE-6, however, the peak speed is increased by a factor of 60. On the other hand, the speed of a single host would be improved only by a factor of 10 or so, if we assume the standard Moore's law (performance doubling time of 18 months). Thus, if we want to achieve a reasonable speed for a similar number of particles as that for GRAPE-4, we need to use around 10 host computers and the communication channel must be 10–20 times faster than that used for GRAPE-4.

Around the time of the design, it was clear that a shared-memory multiprocessor system with 8–16 processors and a sufficient I/O bandwidth would be prohibitively expensive, with a price tag on the order of 1 M USD. On the other hand, a cluster of 8–16 single-processor workstations or PCs would be much less expensive. As far as the cost is concerned, clearly a cluster of single-processor machines would be better than a shared-memory multiprocessor system.

One problem with the cluster is that the simplest configuration (see figure 16) does not work. The reason is as follows.

With this configuration, there are two different ways to distribute particle data over processors (Makino 2002). One is that each processor has a complete copy of the system (the “copy” algorithm). In this case, parallelization is performed as follows. At each blockstep, each processor determines which particles it updates. After all processors update their share of particles, they exchange the updated particles so that all processors have an updated copy of the system. This algorithm has been used to implement the individual timestep algorithm on distributed-memory parallel computers (Spurzem, Baumgardt 1999)

In this algorithm, at the end of the block timestep each processor receives the particles updated on all other processors. This means that the amount of communication is independent of (or, strictly speaking, is a slowly increasing function of) the number of processors, and the overall performance of the system is limited by the speed of communication.

The other possibility is to let each processor to have a non-overlapping subset of the system, so that one particle

resides only in one processor. In this case, with the blockstep algorithm we need to pass around the particles in the current blockstep, so that each processor can calculate the forces from its own particles to particles on other processors (the “ring” algorithm). The amount of communication (host–host and host–GRAPE) per blockstep is again independent of the number of processors. This algorithm is also implemented on distributed-memory parallel computers with direct summation (Dorband et al. 2003) and even with the tree algorithm (Springel et al. 2001).

For general-purpose parallel computers, this simple algorithm actually works rather well, simply because the calculation speed of a single node is so slow. Even a cluster with several hundred nodes is still slower than a single GRAPE-4. Thus, the communication speed of 10–100 MBs⁻¹ is sufficient. However, with GRAPE-6 we do need a faster speed.

We now understand that it is possible to use a hybrid of the above two algorithms to solve the bottleneck (Makino 2002). In this hybrid algorithm, we organize processors into a two-dimensional grid, and distribute the particles so that each row (and each column) has a complete copy of the system.

In the standard realization, this algorithm requires that the total number of processors is r^2 , where r is a positive integer. We divide N particles into r subsets, each with N/r particles. If we number processors from p_{11} to p_{rr} , processor p_{ij} has copies of both the i -th and j -th subsets.

At the beginning of each blockstep, each processor selects the particles to be updated from subset i . Then, all of them calculate the force on them from subset j . After that, the total forces can be calculated by taking a summation over the columns. Here, we assume that the summed results are obtained on diagonal processors p_{ii} .

After particles in the current block are updated on p_{ii} , they are broadcasted to all other processors in the same row (p_{xi}) and also in the same column (p_{ix}), so that both subsets i and j are updated on each processor.

In this algorithm, the amount of communication for one node is $O(N/r)$. In other words, the effective communication bandwidth (both host–host and host–GRAPE) is increased by a factor r . Thus, the communication speed is improved by a factor proportional to the square root of the number of processors.

At present, this solution looks fine, since the price of the fastest single-processor frontend is now rather cheap. The cost of the communication is also rather cheap, with Gigabit Ethernet adapters available for less than 100 USD per unit.

When we started the design of GRAPE-6 in 1996, we did not expect such a drastic change in the price of fast frontend processors. At that time, RISC microprocessors were still several-times faster than PCs with the IA-32 architecture, and 100 Mbit Ethernet adapters were still expensive. Thus, we had to come up with a design that did not need r^2 processors or fast host–host communication.

It was not really difficult to come up with such a design, since the only thing non-diagonal processors do is the force calculation. Instead of a two-dimensional grid of host processors, we can construct a two-dimensional grid of GRAPE hardware with orthogonal broadcast networks (figure 17). The GRAPE hardware in the same row store the same data to their

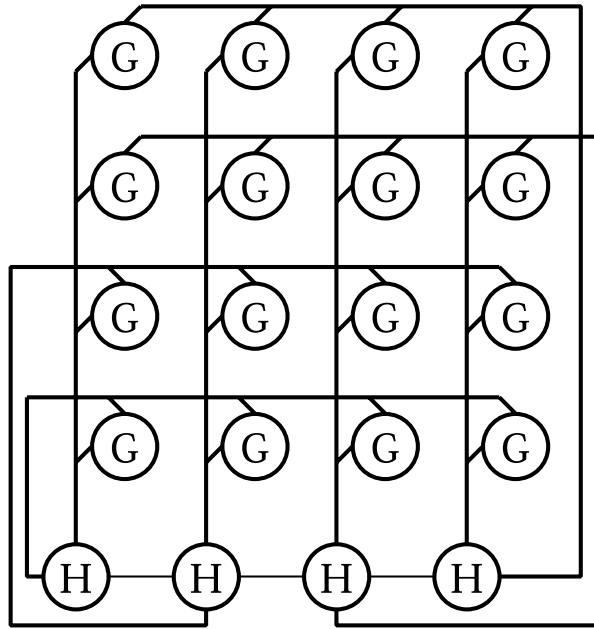


Fig. 17. Two-dimensional network of GRAPE hardware connected to a one-dimensional host network.

particle memories. When they calculate the forces, GRAPEs in the same column receive the same particles and calculate the forces on them from particles in the memory. The calculated results on the boards in the same column are then summed and returned to the host.

One practical problem with this network architecture is that we cannot divide the system to smaller configurations so that we can run multiple programs. In the case of r^2 hosts, we can divide the system to any sub-squares, down to r^2 single host–GRAPE pairs. In the case of a 2D hardware network, we do not have any such division. This problem can be partly circumvented by attaching a simple switching network before the memory interface, so that they can select the input. So we adopted the network structure shown in figure 3.

In the final design of GRAPE-6, we actually adopted a hybrid of the host-grid approach and the GRAPE-network approach, to make a reasonable compromise between the flexibility and absolute performance. Of course, this shift from the pure hardware network to a hybrid one was made partly because we took into account the evolution of the host computers during the development period of GRAPE-6. It has become more cost effective to use a large number of inexpensive (yet fast) computers as the host than to have an elaborate hardware network to connect GRAPEs to a small number of hosts.

5.3. Board–Board Connection

GRAPE-4 consisted of 36 processor boards, organized in a two-stage simple tree network. Nine boards were housed in one rack, with one backplane bus. These boards were all connected to a control board, which broadcasted the data from the host to all processor boards and took the summation of the calculated data on nine processor boards. Since all boards were connected

through a shared backplane bus, the control board had to access processor boards sequentially. In order to improve the data-transfer rate, we used a wide data bus with the width of 96 bits.

The connection between the control board and the host was a 32-bit parallel connection through a coaxial flat cable. This connection was robust and reliable, but had three drawbacks: it was physically large, it was difficult to use long wires, and it was fairly expensive. Because a common clock signal was used on both sides of the connection, the wire length was limited by the allowable signal skew, which meant that it was difficult to use a fast clock (GRAPE-4 used 16 MHz clock).

A more practical problem is that board–board wiring would become too bulky and cumbersome, with hundreds of flat cables and nearly 10000 contact points, if we use the same connection for GRAPE-6. In particular, it would be difficult to design the network board, since it needs to have more than 10 connectors. Also, it would be impractical to use a backplane to connect the network board and processor boards, since the number of pins on the network board would be too large.

An obvious solution to this problem would be to use a fast serial signal, such as the physical layer of the Gigabit Ethernet. At the time of our design decision, however, Gigabit Ethernet was impractical because a copper wire connection was not available in 1998. An optical connection would be too expensive and would dissipate too much heat.

We adopted what is called a “LVDS Link” or a “Flat Panel Display (FPD) link”, which uses four twisted-pair differential signal lines (three for signals and one for the clock). The reason we chose this interface was that inexpensive serializer/deserializer chips were commercially available, and that we could use standard category 5 shielded 4-pair cables for a 100 Mbit Ethernet cable and its connectors for reliable data transmission, for a cable length of up to about 5 meters.

An additional advantage of this choice is that we can use a backplane connection (with custom-designed signal pattern) for connection between the network board and the processor boards. Because the number of signals is small (8 for one port), we can pack many ports into a standard backplane connector (we adopted Compact PCI connector).

5.4. Pipeline Chip and Memory Interface

The processor chip for GRAPE-4 had a single pipeline, which calculated the force on two particles in every six clock cycles (2-way VMP). During the force calculation, the chip received the data of one particle (position, velocity, and mass) in every three external clock cycles, and the width of the input data bus was 107 bits.

One GRAPE-4 board housed 48 pipeline chips, all of which received the same particle data from the memory and calculated the force on two particles. This meant that a single board calculated the forces on 96 particles in parallel.

This shared-memory architecture is simple to implement. However, we could not use this architecture for GRAPE-6, since the hardware parallelism would become excessively large. The pipeline chip for GRAPE-6 would be roughly 50-times faster than that for GRAPE-4. Thus, even if we somehow increase the data transfer rate by a factor of 5, the number of particles on which the forces are calculated in parallel would increase by a factor of 10, from 100 to 1000.

This number is too large if we want to obtain a reasonable performance for simulations of star clusters with small, high-density cores. Note that with multiple-board configurations, this number would become even larger. On an $r \times r$ two-dimensional system, the degree of parallelism becomes larger by a factor of r .

The data-transfer rate of GRAPE-4 chip was about 200 MB s^{-1} . To keep the degree of parallelism to be around 100 or less, the GRAPE-6 chip would have to have a data-transfer rate of 5 GB s^{-1} , which was well beyond our capability of designing and manufacturing. At 100 MHz clock, the speed of 5 GB s^{-1} requires 400 input pins. It is quite difficult to have 400 signal lines, all with a 100 MHz data rate, to connect more than a few chips.

Clearly, a different design was necessary. Excessive degrees of parallelism arose from our decision to let a large number of chips share one memory unit. If we reduce the number of chips to share the memory, we can thus solve the problem. An extreme solution is to attach one memory unit to each pipeline chip, and let multiple pipelines calculate the force on the same set of chips, but from different set of particles.

This extreme solution has one important practical advantage. The connection between the processor chip and its memory is point-to-point, and physically short (since we can put a processor chip and its memory next to each other). This means that a high clock frequency, such as 100 MHz, is relatively easy to achieve.

To attach memory chips directly to the processor chips, we need to integrate the predictor pipeline and the memory controller unit (generation of address and other control signals) to the processor chip. These do not consume many transistors. Therefore, it does not have any effect to the performance of the chip.

With GRAPE-6, we adopted a 72-bit (with ECC) data width for transfer between the memory and the processor chip. A GRAPE-6 chip integrates six 8-way VMP pipelines. Therefore, it calculates the forces on 48 particles in parallel. All pipelines on board calculate the forces on the same set of particles. Thus, even with the largest configuration that we have considered (an 8×8 system), the degree of the parallelism is still less than 400, not much different from that of the full-size GRAPE-4 (which was also 400).

This change from the shared memory design to the local memory design implied that we had to take the summation of a large number of partial forces obtained on chips on one board. With GRAPE-4, we also had to take the summation of forces obtained on different boards, and we used commercially available single-chip floating-point arithmetic units for this summation. With GRAPE-6, we could not apply this solution simply because such chips no longer existed. Thus, we have to either integrate this summation function into the processor chip, or develop another chip to take the summation.

We adopted the latter approach, but used FPGA (field-programmable gate array) chips to implement adders. It was not impossible to integrate floating-point adders into FPGAs, but such a design would require rather large, expensive FPGA chips and a complex design. In order to simplify the design, we chose to use a block floating-point format for the force and the other calculated result. In this format, we specify the

exponent of the result before we start the calculation. Because the actual value of the exponent can be different for forces on different particles, we can calculate the forces with wildly different magnitudes in parallel.

With this block floating-point method, we can greatly simplify the design of the hardware to take the summation. Of course, we have to supply the value of exponent, but its value at the previous timestep is almost always okay. For the initial calculation, we sometimes need to repeat the force calculation a few times until we have a good guess about the exponent.

A rather important advantage of using the block-floating point format is that the calculated result is independent of the number of processor chips used to calculate one force. Since the actual summations, both within the chip and outside the chip, are done in the fixed-point format, no round-off error is generated during summation. Of course, a round-off error is generated when we shift the calculated force to meet the block-floating point format, but this error is independent of the order in which the summation is performed. In the case of the usual floating-point format used in GRAPE-4, the round-off error generated in the summation depends on the order in which the forces from different particles are accumulated, and therefore the calculated force is not exactly the same if the number of boards in the system is different.

Of course, this difference does not have any effect on the accuracy of the simulation, itself, since the word length itself is chosen as such. However, it is quite useful to be able to obtain exactly the same results on machines with different sizes, since it makes the validation of the result much simpler.

6. Performance

In this section, we discuss the performance of the GRAPE-6 system, both for the direct summation algorithm with the individual timestep and the tree algorithm. For both algorithms, we discuss the performance of the single-host system and the multi-host system.

6.1. Direct Summation with Individual Timestep

Here, we discuss the performance of GRAPE-6 for the individual timestep algorithm. As the benchmark runs, we integrate the Plummer model with equal-mass particles for 1 time unit [we use the ‘‘Heggie’’ unit, Heggie and Mathieu (1986), where the gravitational constant G and total mass of the system M are both unity and the total energy of the system E is $-1/4$]. We use the standard Hermite integrator (Makino, Aarseth 1992) with the third-order predictor. The timestep criterion is that of Aarseth (1999) with $\eta = 0.01$. For softening parameter, we tried three different choices. The first one was a constant softening, $\epsilon = 1/64$. We also tried $\epsilon = 1/[8(2N)^{1/3}]$ and $\epsilon = 4/N$ to investigate the effect of the softening size. Note that for $N = 256$, all three choices of the softening gave the same value. In the following, we first describe the performance of a single-host system (with 4 processor boards). We then discuss the performance of a single cluster with 2 or 4 hosts, and finally discuss the performance of multiple-cluster configurations.

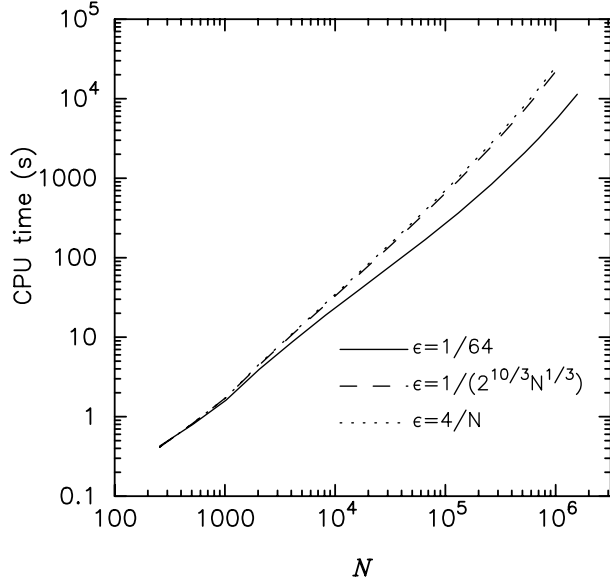


Fig. 18. CPU time in seconds to integrate a Plummer model for 1 time unit plotted versus the number of particles N . The solid, dashed, and dotted curves indicate the result with constant, $1/N^{1/3}$, and $1/N$ softenings, respectively.

6.1.1. Single-host performance

Figure 18 shows the CPU time to integrate the system for one time unit. We actually measured the CPU time for integration from time 0.25 to 1.0 and multiplied the result by 4/3, since the step size after the start of the integration is too small because of the initialization procedure. From figure 18 we can see that the CPU time is almost proportional to N for $N < 10^5$, but for N -dependent softenings the dependence is slightly higher. For $N > 10^5$, the slope approaches 2.

Figure 19 shows the actual calculation speed achieved. The theoretical peak speed of the single-host, 4-PB system is 3.94 Tflop. Here, we define the calculation speed as

$$S = 57 N n_{\text{steps}}, \quad (10)$$

where n_{steps} is the average number of individual steps performed per second. The factor 57 means that we count one pairwise force calculation as 57 floating-point operations. We took this number from recent literatures. From this figure, we can see that the achieved speed is practically independent of the choice of the softening. The reason why calculations with smaller softening take more CPU time is that the number of timesteps is larger, as shown in figure 20. For calculations with N -dependent softenings, the number of block steps increases significantly as we increases N . This means that the average number of particles in one block grows rather slowly. However, as we can see from figure 19, this does not affect the achieved performance.

Roughly speaking, we can model the calculation time per one particle step as follows:

$$T_{\text{single}} = (1 - f)T_{\text{host}} + T_{\text{comm}} + \max(T_{\text{GRAPE}}, fT_{\text{host}}), \quad (11)$$

where T_{host} is the time for the host computer to perform computations to integrate one particle, T_{comm} is the time needed for communication, and T_{GRAPE} is the time to calculate the force on

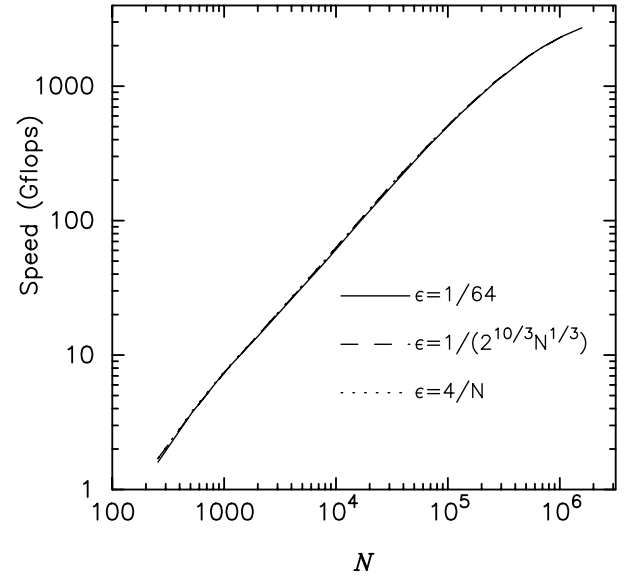


Fig. 19. Same as figure 18, but the calculation speed in Gflops is plotted.

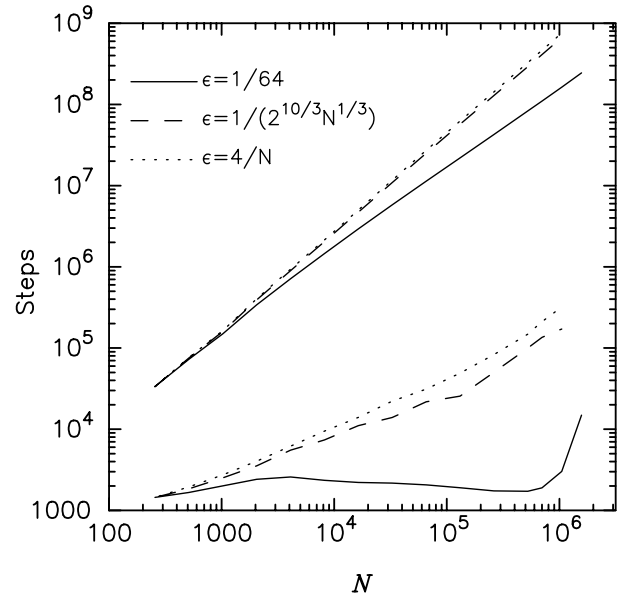


Fig. 20. Same as figure 18, but the number of total individual steps (upper) and block steps (lower) are plotted.

GRAPE. The factor f is the fraction of operations that the host computer can perform while GRAPE is calculating the force. The program we use tries to perform the time integration on host and the force calculation on GRAPE with as much concurrency as possible.

We can estimate T_{comm} as follows. The total amount of data transferred for one particle step is currently 200 bytes. With the present host, the effective data transfer rate for DMA transfer is 80 MB s^{-1} . Therefore,

$$T_{\text{comm}} = 200 / (8 \times 10^7) \text{ s} = 2.5 \times 10^{-6} \text{ s}. \quad (12)$$

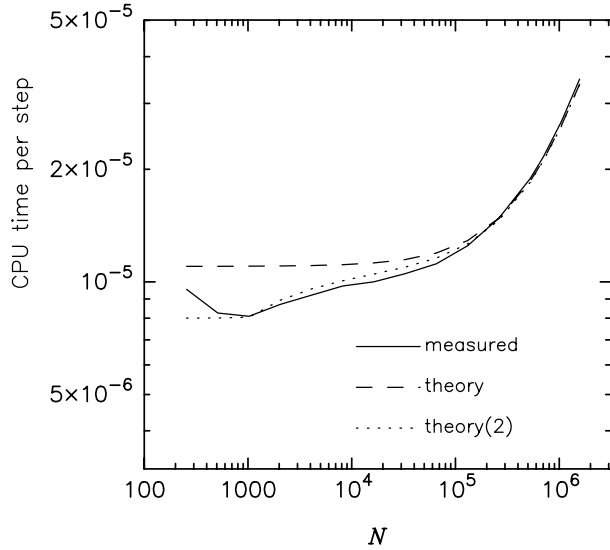


Fig. 21. CPU time per one particle step plotted as a function of the number of particles, N . The solid curve is the measured result. The dashed and dotted curves denote two different theoretical estimates.

The calculation time on GRAPE is expressed as

$$T_{\text{GRAPE}} = N / (9 \times 10^7 n_{\text{pipes}}) \text{ s} = 1.447 \times 10^{-11} N \text{ s}, \quad (13)$$

where n_{pipes} is the total number of pipelines. With our current system $n_{\text{pipes}} = 768$.

In figure 21, the solid curve shows the measured CPU time per step. The dashed curve is a fit, with $T_{\text{host}} = 8.5 \times 10^{-6}$ s and $f = 0$. We can see that agreement between the theory and experimental result is good for large N , but is rather poor for small N . This is because we ignored the effect of the cache memory on T_{host} . The dotted curve is the theoretical estimate with a heuristic model for the cache effect. For this curve, we used

$$T_{\text{host}} = 5.5 \times 10^{-6} c + 8.5 \times 10^{-6} (1 - c) \text{ s}, \quad (14)$$

where c is expressed as

$$c = \begin{cases} 1, & (N \leq 1000), \\ \sqrt{N/1000}, & (N > 1000). \end{cases} \quad (15)$$

This model is purely empirical, but apparently gives a reasonable description for the performance. Since this effect of the cache is rather large, it turned out to be difficult to determine the value of f empirically. We assumed $f = 0$.

For $N < 1000$, the experimental value is larger than the prediction of the refined theory. This is because the number of particles in one block is too small. The overhead to invoke DMA operations becomes visible.

Up to here, we have discussed only the speed of a 4-PB system. Since there are many installations of GRAPE-6 outside Tokyo university with one PB connected to a host, it would be useful to give the performances of smaller configurations. Figure 22 gives the estimated performance of the 4-, 2-, and 1-PB system. One can see that performance difference is rather small for $N < 3 \times 10^4$. For $N > 10^5$, the performance difference becomes significant.

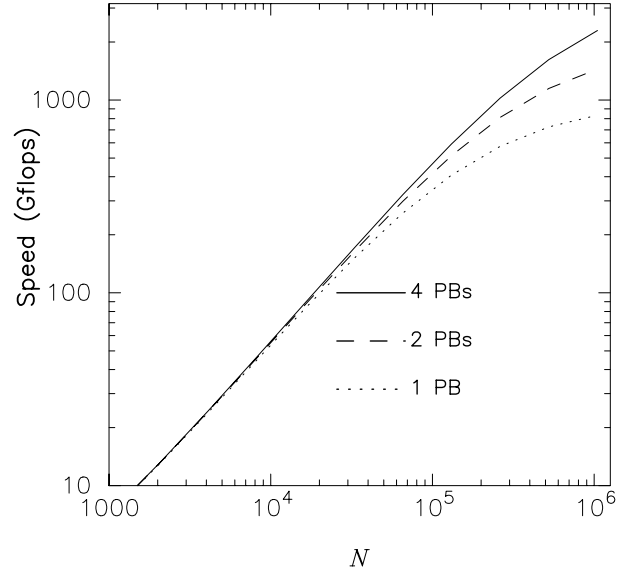


Fig. 22. Estimated performance of the 4-, 2-, and 1-PB systems as a function of the number of particles, N . The solid, dashed and dotted curves denote the speed of the 4-, 2-, and 1-PB systems, respectively.

6.1.2. Multi-host performance

Figure 23 shows the calculation speed for multi-host systems with up to 4 hosts. The peak speeds of the 2- and 4-host systems are 7.88 Tflops and 15.76 Tflops, respectively. For up to 4 hosts, the network boards are used to distribute the data, and the communication network between the host computers are used primarily for synchronization. The parallel program, itself, was written using MPI, and we used MPICH/p4 over TCP/IP as the MPI library. The network interface is Planex GN-1000TC Gigabit NIC, which uses the NS 83820 chip. We found the performance of MPICH/p4 on this network interface to be quite unsatisfactory, and used UNIX TCP/IP socket system calls for actual communication.

We can see that multi-host codes require a rather large number of particles to achieve a speed faster than that of the single-host code. Even with constant softening, the two-host code becomes faster than the single-host code only at $N \sim 3000$, and for $\epsilon = 4/N$ this crossover point moves to around $N \sim 10^4$.

Figure 24 shows the calculation time per one particle step for a 4-node parallel calculation. The measured value is obtained by dividing the total number of particle steps by the wallclock time. This figure clearly shows why the value of N for the crossover is rather large. For a “small” N ($N < 10^4$), the calculation time is inversely proportional to the number of particles, N . This is because the communication between hosts, which takes constant time per one blockstep, dominates the total cost in this regime. To be quantitative, the calculation time per one particle step is expressed as

$$T_{\text{mn},p} = T_{\text{single}}/p + T_{\text{comm,hosts}}, \quad (16)$$

where p is the number of nodes in a cluster and $T_{\text{comm,hosts}}$ is the communication time, expressed as

$$T_{\text{comm,hosts}} = 6(\log_2 p + 1)t_{\text{sync}}/n_b, \quad (17)$$

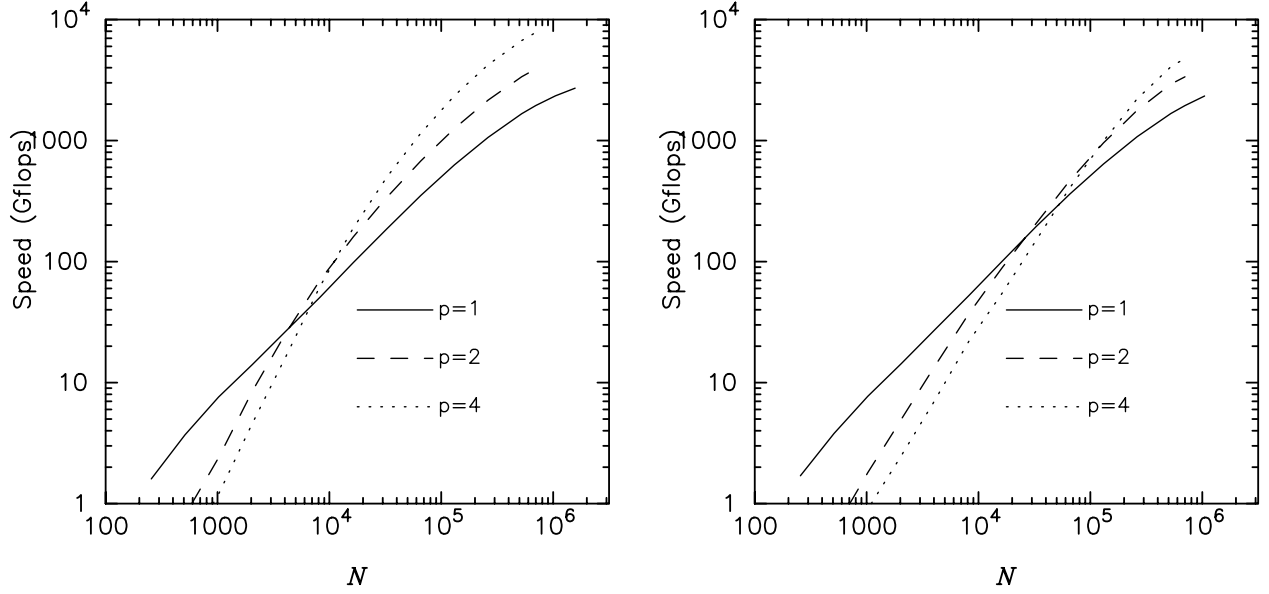


Fig. 23. Calculation speed in Gflops plotted as a function of N . The solid, dashed, and dotted curves show the results for the 1, 2, and 4-node systems, respectively. The left panel shows the result for constant softening, and the right panel for $\epsilon = 4/N$.

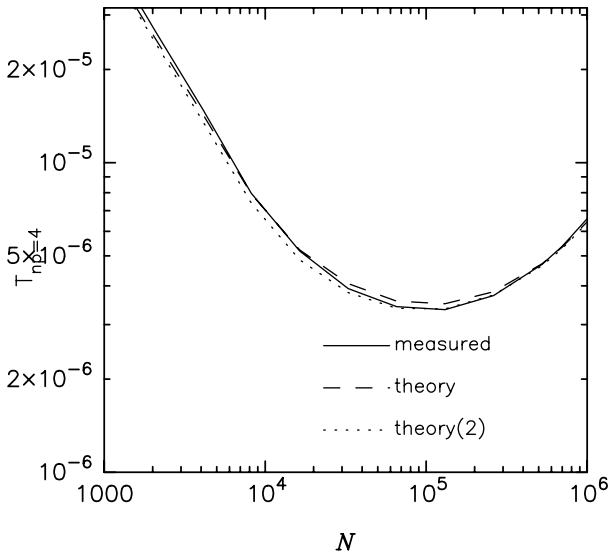


Fig. 24. Same as figure 21, but for the case of a 4-node parallel calculation.

where $(\log_2 p + 1)t_{\text{sync}}$ is the time to complete a barrier synchronization for the parallel code running on p nodes. The logarithmic factor comes from the fact that synchronization requires $\log_2 p + 1$ stages. The divisor, n_b , is the average number of particles integrated in one blockstep. For our current implementation of the synchronization, we found $t_{\text{sync}} = 250 \mu\text{s}$. The factor 6 is the number of synchronization operations necessary in one blockstep.

The theoretical estimates shown in figure 24 were calculated using equation (16). Here again, the agreement between the measured result and the theory with the effect of the cache memory of the host is very good. To evaluate T_{host} , we used

N/p instead of N in equation (15), since one node handles N/p particles.

6.1.3. Multi-cluster performance

Figure 25 shows the calculation speed for multiple-cluster systems, as a function of the number of particles in the system, N . The crossover point at which multi-cluster systems becomes faster than single-cluster system is rather high ($N \sim 10^5$), and even for $N = 10^6$, the speedup factors achieved by multi-cluster systems are significantly smaller than those for the ideal speedup.

For a multi-cluster system, the calculation time per one particle step can be estimated as follows. In our current implementation of the multi-cluster calculation code, one host of a p -hosts, q -cluster system (therefore p/q hosts in a cluster) handles N/p particles. The forces from the particles in the hosts in the same cluster can be calculated using the hardware network on the side of GRAPE-6. However, one cluster needs to gather information on particles of different clusters. By letting each of the p/q hosts in one cluster receive data from other $q - 1$ hosts, we can let one cluster maintain complete date of all N particles. This is just one of many possible implementations. For a small value of q , theoretically, this is close to the best possible implementation.

With this implementation, the calculation time per one particle step is expressed as

$$T_{\text{mc},p} = T_{\text{single}}/p + T_{\text{comm,hosts}} + T_{\text{comm,clusters}}, \quad (18)$$

where $T_{\text{comm,clusters}}$ is the time for communication between hosts in different clusters. It is expressed as

$$T_{\text{comm,clusters}} = 72(2t_{\text{comm,net}} + t_{\text{comm,grape}})(q - 1)/p, \quad (19)$$

where $t_{\text{comm,net}}$ and $t_{\text{comm,grape}}$ are the time to send 1-byte date through the network interface of the host and the host-GRAPE interface, respectively. The constant factor of 72 is the length of

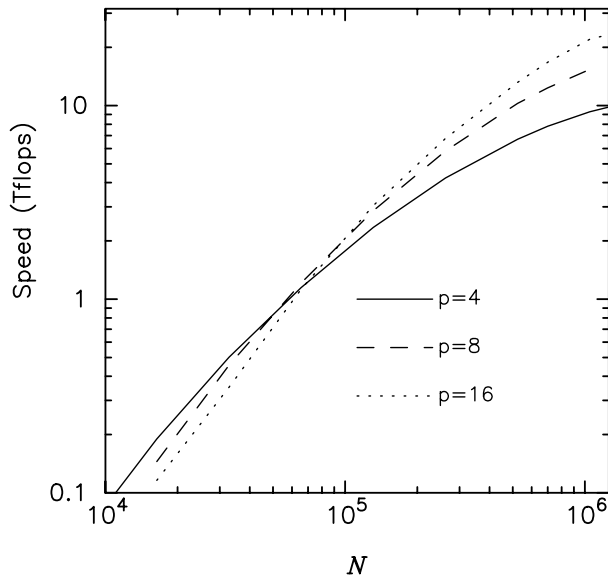


Fig. 25. Calculation speed in Tflops plotted as a function of N . The solid, dashed, and dotted curves show the results for the 4, 8, and 16-host (1, 2, and 4-cluster) systems, respectively. Constant softening is used for all runs.

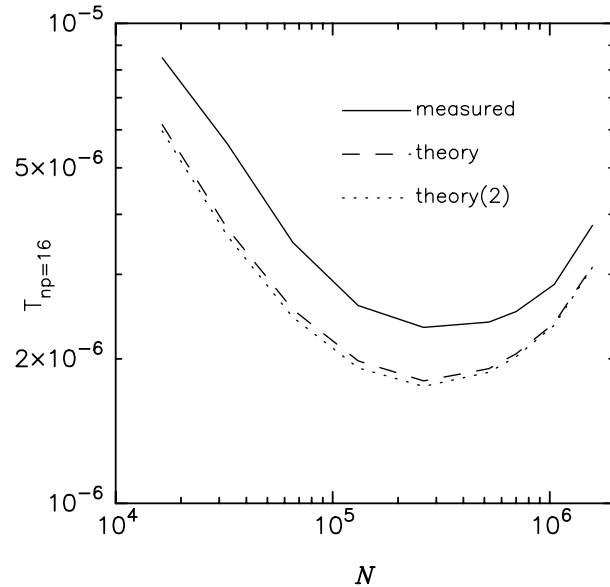


Fig. 26. Same as figure 21, but for the case of a 16-host parallel calculation.

data for one particle in bytes. The next factor of 2 comes from the fact that each node needs to both send and receive data. The factor $q - 1$ appears since one node receives data from $q - 1$ other nodes. We used $t_{\text{comm,net}} = 1.7 \times 10^{-8}$ s and $t_{\text{comm,grape}} = 1.25 \times 10^{-8}$ s. These values are based on separate measurements using small benchmark programs. Figure 26 shows the calculation time per one particle step for a full-cluster calculation (16 nodes, 4 clusters). The agreement between the theoretical estimate and the measured value is fairly good, but not ideal. We probably have underestimated $t_{\text{comm,net}}$ in the real program.

6.1.4. Summary for the direct summation code

In this section we present the performance figures of GRAPE-6 for a direct N -body simulation. As described in the introduction, this kind of simulations is the main target of GRAPE-6. This, we made a fairly detailed analysis of the performance. What we have found is summarized as follows.

In the case of the single-host configuration, the calculation speed of the present host computer is the largest bottleneck of the performance, and the communication speed is relatively unimportant. This means that we can keep improving the overall performance of the system just by replacing the host computer, for the next several years.

For the multi-host configuration, the situation is rather different. In the case of a single cluster (no host–host data transfer), the performance for small- N runs is determined by the overhead of the barrier synchronization between host computers. We currently use standard UNIX implementation of the TCP/IP socket for the basic communication; the TCP/IP socket is certainly not the communication software with the lowest possible latency. The use of communication software/hardware with a lower latency would significantly improve the performance.

Finally, for the case of the multi-cluster configuration, as expected, the performance is limited by the bandwidth of the communication between the hosts. Currently, we use a Gigabit Ethernet card on a 32-bit, 33-MHz PCI bus. Clearly, by going to a faster bus (PCI-66 or PCI-X) and a faster CPU, the communication bandwidth will be significantly improved.

To summarize, at present the performance of GRAPE-6 for small- N calculations is limited by the speed of the host computer, and by the latency of the communication between the hosts when multi-host or multi-cluster systems are used. Even so, GRAPE-6 can achieve a speed exceeding 100 Gflops, for relatively small number of particles, such as 16k. In the coming several years, improvements of the host computer will improve the overall performance of the system.

6.2. Tree Algorithm

Here, we discuss the performance of GRAPE-6 for the Barnes–Hut tree algorithm. We use the modified algorithm introduced by Barnes (1990). We discuss the performance of a single-host code and a multi-host (parallel) code with up to 12 host computers. The parallel algorithm is based on a space decomposition similar to the well-known orthogonal recursive bisection (ORB) method (Dubinski 1996). The details of the parallel algorithm will be discussed elsewhere.

6.2.1. Single-host performance

Figure 27 shows the CPU time per timestep as a function of the number of particle, N . The distribution of particles is a Plummer model, with an outer cutoff radius of 22.8 in Heggie units. We used $n_g = 20000$ as the maximum group size for the modified algorithm.

We can see that the CPU time grows practically linearly as we increase N . Also, the dependence on the opening angle, θ , is rather weak. This weak dependence is the characteristic of GRAPE implementation of the tree algorithm (Makino 1991c;

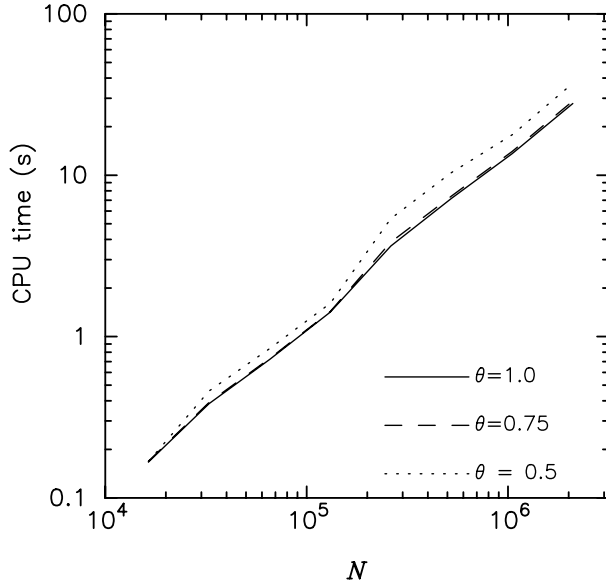


Fig. 27. CPU time in seconds per timestep plotted as a function of the number of particles, N , on a single-host configuration. The solid, dashed, and dotted curves are for $\theta = 1, 0.75, 0.5$, respectively.

Table 4. CPU time distribution for the treecode.

Operation	Time (s, $\theta = 1$)	($\theta = 0.5$)
Tree construction	7.57	7.57
Force calculation	18.40	27.62
Other operations	1.86	1.86
Total	27.83	37.05

Athanassoula et al. 1998).

Table 4 gives a breakdown of the CPU time per step for calculations with $N = 2^{21}$. The average length of the interaction list is 1.01×10^4 and 1.69×10^4 for $\theta = 1.0$ and 0.5 , respectively. The number of groups is 310 for both cases. As in the case of a tree algorithm on older GRAPE hardware, the performance is limited by the speed of the host and that of communication. The actual calculation on GRAPE-6 takes less than three seconds for the case of $\theta = 0.5$. The calculation on the host (tree construction, tree traversal, and other calculations including the data conversion between GRAPE-6 format and floating-point format) counts for roughly 2/3 of the remaining time, and actual communication 1/3. This, again, implies that there is a rather large room for the improvement of the speed, just by moving to faster host computers.

6.2.2. Multi-host performance

Since the performance of the single-host GRAPE-6 is limited by the speed of the host computer, an obvious way to improve the performance is to use multi-host systems.

Figure 28 shows the performance of the parallel tree algorithm. The program used is a newly written one based on orthogonal recursive multi-section, a generalization of the widely used ORB tree that allows a division to an arbitrary

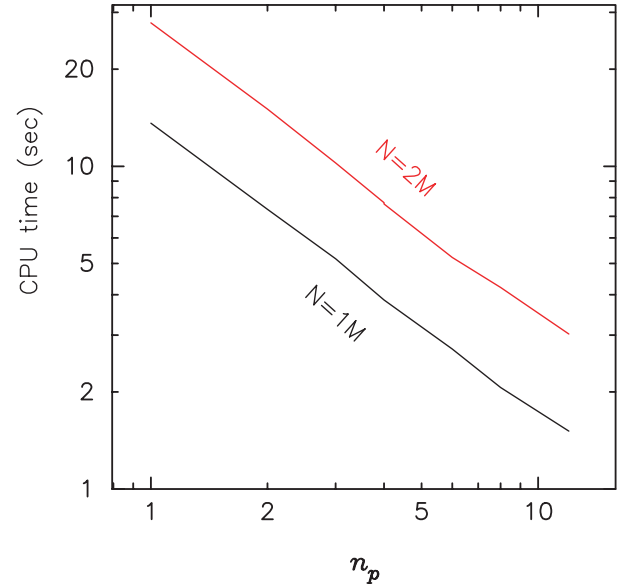


Fig. 28. CPU time in seconds per timestep plotted as function of the number of host computers, n_p . The parallel tree algorithm is used with $\theta = 1.0$. The upper and lower curves are the result for 2^{21} and 2^{20} particles, respectively. The initial distribution of the particles is the Plummer model.

number of domains in one dimension, instead of allowing only bisection. The primary advantage of this algorithm is that it can be used when the number of host computers is not exactly a power of two. We measured the performance on 1, 2, 3, 4, 6, 8, and 12 hosts.

The distribution of the particles is again the Plummer model. One can see that the scaling is again quite good. A 12-host calculation is 9.3 times faster than a single-host calculation. The parallel efficiency is better than 75%, even for the relatively small number of particles shown here.

7. Discussion

7.1. Hintsights

Though we regard GRAPE-6 to be a reasonable success, this certainly does not mean that we have done everything correctly. We did make quite a few mistakes, some of which affected the performance, some affected the reliability, some extended the development time, and some limited the application range. In the following we briefly discuss them in turn.

7.1.1. Performance

Concerning the performance, the largest problem with GRAPE-6 is that its clock frequency is somewhat below the expected value. The design goal (for the “worst case”) was 100 MHz, while our actual hardware is currently running at 90 MHz. With GRAPE-4, the design goal was 33 MHz, and the machine operated without any problem at 32 MHz. The processor chip itself was confirmed to operate fine at 41 MHz.

The primary reason for the low operating frequency is a problem with the stability of the power supply to the chip, or the impedance of the power line. Compared to the GRAPE-4

processor chip, the GRAPE-6 processor chip consumes about a two-times larger power at half the supply voltage. Thus, to keep the relative drop of the supply voltage to be the same, the impedance of the power line must be $1/8$ of that of GRAPE-4.

This is a quite difficult goal to achieve. Initially, even the manufacturer of the chip did not fully appreciate how hard it would be. As a result, the first sample of the chip could not operate correctly at a clock speed higher than 60 MHz. The problem was that, when calculation starts, the power consumption of the chip increases by about a factor of two compared to that when the chip is idle. Because of the large total resistance of the power line in the LSI package and the power plane of the silicon chip, itself, the supply voltage to the transistors decreases, and as a result their switching speed slows down.

In the second design, the manufacturer came up with an additional power plane and an increased number of power and ground pins, which reduced the resistance significantly. However, the result was still rather unsatisfactory.

The manufacturer was not alone in making this kind of mistake. In the design of the processor board and the power supply, we also made similar mistakes. In the first design, we used traditional large switching power units with relatively low switching frequency. This unit turned out to be unable to react to a quick change in the load between the idle and calculation states. The normal electrolytic capacitors also turned out to be completely useless in stabilizing the power supply voltage. Thus, we need to redesign the power supply unit with high-frequency inverters and low-ESR capacitors.

In hindsight, we could have borrowed the design of power supply units for standard PC motherboards (for Intel processors), which were designed to meet quite similar requirements, but for an extremely low cost. The power supply circuit for a typical PC motherboard would be good enough to support single module with 4 chips. We could then supply 12 V to PCB.

These apparently minor technical details are absolutely crucial for manufacturing of high-performance computers.

Another problem with the current GRAPE-6 chip is its limited I/O performance of only 90 MB s^{-1} . As we stated earlier, this bandwidth is sufficient to keep the standard PCI interface busy, and it is not really a bottleneck, since, for many applications, calculations on the host computer are more time-consuming. Even so, in a few years the I/O performance will become a problem. An additional problem is that host computers with faster PCI interfaces (PCI64 and PCI-X) are now available. We cannot take advantage of these faster interfaces with the current GRAPE-6 design, because the I/O bandwidth of the processor chip is limited. We could have increased the I/O bandwidth of the chip without too much problem, by allowing a change in the ratio between the chip clock and board clock. With our current design, this ratio is effectively fixed at 4.

Even with the current chip design, we could have increased the communication bandwidth of the processor board without increasing that of the chip, by letting multiple chips transfer the data simultaneously. This possibility should have been considered to increase the lifetime of the hardware.

7.1.2. Reliability

Since the GRAPE-6 system consists of an exceptionally large number of arithmetic units, one might imagine that the primary source of the error is the calculation logic, itself. In practice, however, we have almost never seen any calculation error, once the power supply had become good enough. On the other hand, we found quite a few errors in data transfer.

We implemented an ECC circuit for the memory interface of the processor chip, but only added a parity detection circuit to the I/O ports. We thought that this is reasonable, since the memory ports operate on a 90 MHz clock and the I/O ports on 22.5 MHz. However, it turned out that a memory parity error almost never occurs, while parity errors for I/O occur rather frequently. Since we do not exactly know the type of error, it is not 100% clear whether the ECC capability would have helped or not. However, it is at least clear that more reliable data transmission would be better.

A more serious problem concerning the reliability was a very high defect rate for mass-produced processor boards and processor modules. Practically all failures were due to unreliable soldering, and most of the soldering problems turned out to be simply due to a lack of skill of the manufacturer. This may be telling something about the present performance of the Japanese high-tech industry. Even so, it is certainly true that to manufacture a rather small quantity of PCB boards is difficult. We could have designed either more automated test procedures for boards (with JTAG standard) or made redundant connections. Yet another possibility was to reduce the number of wires by using higher-frequency signals.

7.1.3. Development time

As discussed in section 5, the use of a parallel host was inevitable. However, the use of a multicast network was not, at least in hindsight. We assumed that the price of high-end uniprocessor computers would not change greatly, and that the cost of high-bandwidth network adapters (1 Gb s^{-1} or higher) would remain high. In other words, we assumed that we could not afford to buy ~ 100 fast host computers and to connect all of them by a fast network. Therefore, we designed our own network, which connects r host computers to r^2 processor boards. This approach worked fine, as we have shown in the previous section. However, an alternative design, in which we connect each processor board to its own host computer, would have been much easier to develop.

In 1997, the fastest systems were RISC-based UNIX workstations, with price higher than 20k USD. In 2003, systems based on the Intel x86 architecture offer a speed similar to that of the RISC-based systems with the highest performance for a cost of less than 2k USD. To illustrate this, we use SPECfp (either 95 or 2000) numbers as being representative of the performance. In 1997 the speed difference between the RISC systems and the x86 systems was nearly a factor of three. This ratio was almost constant during 1990s. The reason why the ratio started to shrink is simply that the rate of improvement of the performance of RISC-based systems slowed down. Thus, it would have been difficult to predict the present state in 1997, or even in 1999. In other words, even though it is now clear that the network hardware of GRAPE-6 is not necessary, until 2000 we had no other choice.

Another reason for the rather long development time (aside from the problems with the power supply) is the fact that we effectively integrated all functions of the system to the processor chip. It integrates the memory controller, the predictor pipeline, and all other control logics. Except for the predictor pipeline, all of these logics were implemented with FPGAs on GRAPE-4. This integration simplified the design of the board, but fortunately we have not made any serious mistake in the design of these parts. The integration of these complicated logics onto hardware required extremely careful design and test procedures, which were time-consuming. With the present price of moderately large FPGA chips, all of these control logics could be implemented using FPGAs, with very small additional cost. Of course, such a moderately large and inexpensive FPGA was not available when we decided on the design of the GRAPE-6 chip. However, we could have predicted the direction of the evolution of FPGA chips and estimated their price.

7.1.4. Application range

Since GRAPE-6 is designed solely for the gravitational N -body problem, one might think that there is not much range of applications. However, even within N -body simulations, there are many factors.

The overall design of GRAPE-6 is highly optimized to parallel execution of the direct force calculation with the individual timestep algorithm. This of course means that it is not optimal for other applications, such as the tree algorithm and SPH calculations of a self-gravitating fluid.

With the case of the tree algorithm, the performance is limited mainly by the speed of the host computer. Thus, in this case, adding more host computers would have greatly improved the performance.

In principle, we could have improved the performance of the tree algorithm in several other ways. One obvious approach is to reduce the data to send. With the tree algorithm, we would not use a predictor. Moreover, we would not need the full 64-bit resolution for the position data. Thus, we could have implemented some way to reduce the data to send for j -particles, if our memory controller was not implemented in hardware. Actually, the memory controller of GRAPE-6 has some programmability. However, one “feature” of this memory controller prevented us from taking full advantage of this programmability to reduce the amount of data transfer.

With an FPGA implementation of the memory controller, we could implement other ways to further reduce the communication. For example, we could implement indirect addressing, so that we could send indices of j particles instead of sending their physical data.

Concerning the design of the pipeline, one thing which might have been useful for simulating collisionless systems or composite N -body+SPH systems is the ability to apply different softening lengths on different particles in a symmetrized way. This can be achieved by calculating the softened distance as

$$r_s^2 = r_{ij}^2 + \epsilon_i^2 + \epsilon_j^2. \quad (20)$$

The pipeline will need one more addition, which is relatively inexpensive.

With SPH, the main problem is that the calculation of SPH interactions, itself, cannot be done on GRAPE-6. The PROGRAPE system (Hamada et al. 2000), with the calculation pipeline fully implemented in FPGA, could be used to calculate SPH interaction. An moderately large PROGRAPE system is currently under development.

With the logic design of the pipeline, we have noticed a few problems which we could not foresee. One is the length of the accumulator for the time derivative of the force. For the force and potential, we used 64-bit accumulators, but for the time derivative we used 32-bit accumulators. As far as the accuracy is concerned, this length is long enough. However, when we performed simulations with a large number of particles, we realized that overflow occurred rather frequently. The reason why overflows occurs is that the magnitude of the time derivative of the force can change by a large factor in a single timestep. A large change occurs when the previous value happens to be almost zero. We could circumvent this problem by combing a guess for the likely value of the time derivative of the force based on the value of force and timestep, but it is cumbersome to implement and expensive to evaluate. By increasing the accumulator length to, say, 40 bits, we could have almost completely eliminated the overflow. This overflow does not have any noticeable impact on the performance. However, the need to handle overflows made the interface program rather complicated.

7.2. GRAPE-7/8

Given that GRAPE-6 is now completed and we have already had the experience of running it for almost two years, it would be natural to put some thought into how its successor will look. In this section, we first discuss changes in technologies, and then give an overview of the design possibilities.

7.2.1. Technological changes and the basic design

Compared to the technology used in GRAPE-6, what will be used in the next GRAPE system (we call it NGS for short) will be different in:

- (a) semiconductor technology,
- (b) development cost,
- (c) host I/O bus.

First, let us discuss semiconductor technology. GRAPE-6 uses $0.25\mu\text{m}$ technology, while NGS would use, depending on the time to start, either 130 nm or 90 nm technology. Since it seems that we will not get the necessary budget very soon, we will probably use 90 nm. This means that we can pack about 8-times more transistors into a chip of the same size, and the switching speed will be about 3-times faster. Thus, a single chip of the same size can offer 20-times more computing power. If the power supply voltage is reduced by the same factor, the power consumption would remain the same, but most likely the supply voltage would be somewhat higher, resulting in a significant increase in the power consumption.

To express in concrete numbers, a single chip would integrate around 50 pipeline processors, each with 60 arithmetic units operating at a 300 MHz clock speed, with a 1.2 V supply voltage and power consumption of 20 W. The theoretical peak speed of the chip will be around 600 Gflops.

Compared to the projected speed of general-purpose microprocessors in, say, 2007, this speed is quite attractive. In 2007, microprocessors will, at best, have a peak speed 10-times faster than they have now, or about 30 Gflops. The typical performance on real applications would be around 10 Gflops or less, for a power consumption of 100 W or more.

A necessary consideration is how to connect the pipelines to memory. If we use the same memory system as we used for GRAPE-6, the total number of virtual pipelines per chip becomes 1000, which is too large for a simulation of any collisional system. As is the case with GRAPE-6, it is necessary to keep the number of i -particles calculated in parallel to be around 500 or less for large systems with many chips. Thus, the number of virtual pipelines per chip must be less than 200, or ideally less than 100. In other words, the memory bandwidth must be increased by at least a factor of five, to around 3.5 GB s^{-1} .

This number, by itself, sounds relatively easy to achieve. It is the same as what was used with the first Intel P4 processor (3.2 GB s^{-1}), using two DRDRAM channels each with a 16-bit data width. The Intel P4 has been around for more than two years. We can now also use DDR 400 memory chips, which have 4-times more throughput than the SSRAM chips used in GRAPE-6. We could also use DDR SRAMs.

The choice of the memory interface has strong impact on the range of applications. One major limitation of GRAPE-6 is that, as discussed in the previous section, its memory addressing scheme is limited only to sequential access to a full set of predictor data. Thus, it is not easy to use the tree or other sophisticated algorithms efficiently on GRAPE-6. One possibility to solve this problem is to implement the memory controller and other control logics in an FPGA chip. The connection between the FPGA chip and the pipeline chip must be quite fast, but this is relatively easy to achieve, since the data transfer is unidirectional from the FPGA chip to the pipeline chip. The memory controller will be implemented in the FPGA chip. Thus, it will be possible to use different types of memory (DRDRAM, DDR DRAM/SRAM) without any need to change the pipeline chip.

As we discussed earlier, parallelism will be achieved by a two-dimensional network of host computers. Each of them will have a relatively small GRAPE system. As an example, we consider a system with 256 host computers, each with two GRAPE cards. Each card houses 4 processor chips with their own memory control units and memories. All of them can be packaged into a single card of the PCI form factor, though we need to use special care for the power supply.

For the interface to the host, the easiest solution is to use PCI-X, which is now available with a data transfer speed of up to 1 GB s^{-1} . PCI-X gives us an order-of-magnitude increase in the communication speed, which roughly balances the increase in the performance of a factor of 20. One problem is whether or not PCI-X will be available around 5 years from now. We need to predict the market trend, or develop a design that can use multiple interfaces.

Note that this factor-of-10 increase in the communication implies that the chip-to-chip communication must also be faster by the same factor. This is not easy, but since the physical size of the board will be much smaller, it would not

be impossible to use fast clocks.

Thus, the design of NGS seems to be simple, as long as we set the parallel execution of the individual timestep algorithm with direct summation as the primary design target.

The only, but quite serious, problem is that the predicted initial cost for the custom chip will be very high. The initial cost for a custom chip has been increasing quite steeply. Roughly speaking, the initial cost has been proportional to the inverse of the design rule. Thus, while the initial cost of the GRAPE-4 chip was around 200 k USD, that for GRAPE-6 exceeded 1 M, and for NGS it will reach 3 M. Even though this is “small” compared to the price of any massively-parallel supercomputer or even PC clusters, to obtain a grant of this size within the small community of theoretical astrophysics in Japan is not easy.

7.2.2. *Combination with sophisticated algorithms*

One rather fundamental question concerning the next GRAPE system is whether direct summation is really the best solution or not. McMillan and Aarseth (1993) have demonstrated that it is possible to implement a combination of the Barnes–Hut tree algorithm and the individual timestep algorithm that runs efficiently at least on a single-processor computer, and potentially also on shared-memory parallel computers. Even when we require very high accuracy, the gain by the tree algorithm is large for a large N . For example, the number of interactions per particle to achieve a relative force accuracy of 10^{-5} is around 8000 when the quadrupole moment is used, and around 2000 when the octupole moment is used, for around 10^6 particles. Thus, even if we assume that the calculation cost of an octupole is 10-times higher than that of a point-mass force, the calculation cost of the tree algorithm would be a factor of 50 less than that of a direct calculation.

Even though the scaling is not as drastic as that of the tree algorithm, the Ahmad–Cohen scheme (1973, also known as the neighbor scheme) offers a quite significant reduction of the calculation cost over the simple direct summation. The theoretical gain in the calculation cost is $O(N^{1/4})$ for the neighbor scheme (Makino, Hut 1988; Makino, Aarseth 1992). However, the actual speedup is nearly a factor of 10 for only 1000 particles. Thus, for 10^6 particles the gain can reach a factor of 50.

For 10^6 particles, both the tree algorithm and neighbor scheme, at least theoretically, offer reduction in the calculation cost of around a factor of 50. This factor is certainly still smaller than the advantage of the GRAPE hardware over general-purpose computers, since the difference in the price–performance ratio will exceed 10^3 . However, if we can incorporate either of these sophisticated algorithms, even with a significant loss in the hardware efficiency, like a factor of 5 or even more, we can still achieve a very significant improvement in the overall speed. We are currently investigating several possible ways to achieve this goal.

We would like to thank all of those who are and have been involved in the GRAPE project. In particular, we thank Daiichiro Sugimoto for his continuous support to the project, Atsushi Kawai for helping with the hardware design, Yoko Funato, Eiichiro Kokubo, Simon Portegies Zwart, Piet Hut, Steve McMillan, Makoto Taiji, Sverre Aarseth, Takayuki Saito

and many others for discussions on their experience with GRAPE-4 and 5. We are grateful to Mary Inaba and Junichiro Shitami for discussions on the network performance, and for providing us with their software to measure the network

performance. This work is supported by the Research for the Future Program of Japan Society for the Promotion of Science (JSPS-RFTF97P01102).

References

- Aarseth, S. J. 1963, *MNRAS*, 126, 223
 Aarseth, S. J. 1999, *Celest. Mech. Dyn. Astron.*, 73, 127
 Ahmad, A., & Cohen, L. 1973, *J. Comput. Phys.*, 12, 389
 Athanassoula, E., Bosma, A., Lambert, J. C., & Makino, J. 1998, *MNRAS*, 293, 369
 Barnes, J. E. 1990, *J. Comput. Phys.*, 87, 161
 Barnes, J., & Hut, P. 1986, *Nature*, 324, 446
 Dorband, E. N., Hensendorf, M., & Merritt, D. 2003, *J. Comput. Phys.*, 185, 484
 Dubinski, J. 1996, *New Astron.*, 1, 133
 Fukushige, T., Ito, T., Makino, J., Ebisuzaki, T., Sugimoto, D., & Umemura, M. 1991, *PASJ*, 43, 841
 Fukushige, T., Taiji, M., Makino, J., Ebisuzaki, T., & Sugimoto, D. 1996, *ApJ*, 468, 51
 Funato, Y., Makino, J., & Ebisuzaki, T. 1992, *PASJ*, 44, 613
 Greengard, L., & Rokhlin, V. 1987, *J. Comput. Phys.*, 73, 325
 Hamada, T., Fukushige, T., Kawai, A., & Makino, J. 2000, *PASJ*, 52, 943
 Heggie, D. C., & Mathieu, R. D. 1986, in *The Use of Supercomputers in Stellar Dynamics*, ed. P. Hut & S. McMillan (New York: Springer), 233
 Ito, T., Ebisuzaki, T., Makino, J., & Sugimoto, D. 1991, *PASJ*, 43, 547
 Ito, T., Makino, J., Ebisuzaki, T., & Sugimoto, D. 1990, *Comput. Phys. Comm.*, 60, 187
 Kawai, A., Fukushige, T., Makino, J., & Taiji, M. 2000, *PASJ*, 52, 659
 Makino, J. 1991a, *PASJ*, 43, 859
 Makino, J. 1991b, *ApJ*, 369, 200
 Makino, J. 1991c, *PASJ*, 43, 621
 Makino, J. 2002, *New Astron.*, 7, 373
 Makino, J., & Aarseth, S. J. 1992, *PASJ*, 44, 141
 Makino, J., & Hut, P. 1988, *ApJS*, 68, 833
 Makino, J., Kokubo, E., & Taiji, M. 1993, *PASJ*, 45, 349
 Makino, J., & Taiji, M. 1998, *Scientific Simulations with Special-Purpose Computers — The GRAPE Systems* (Chichester: John Wiley and Sons)
 Makino, J., Taiji, M., Ebisuzaki, T., & Sugimoto, D. 1997, *ApJ*, 480, 432
 McMillan, S. L. W. 1986, in *The Use of Supercomputers in Stellar Dynamics*, ed. P. Hut & S. McMillan (New York: Springer), 156
 McMillan, S. L. W., & Aarseth, S. J. 1993, *ApJ*, 414, 200
 Okumura, S. K., et al. 1993, *PASJ*, 45, 329
 Okumura, S. K., Ebisuzaki, T., & Makino, J. 1991, *PASJ*, 43, 781
 Springel, V., Yoshida, N., & White, S. D. 2001, *New Astron.*, 6, 79
 Spurzem, R., & Baumgardt, H. 1999, *MNRAS* submitted
 Steinmetz, M. 1996, *MNRAS*, 278, 1005
 Sugimoto, D., Chikada, Y., Makino, J., Ito, T., Ebisuzaki, T., & Umemura, M. 1990, *Nature*, 345, 33
 Umemura, M., Fukushige, T., Makino, J., Ebisuzaki, T., Sugimoto, D., Turner, E. L., & Loeb, A. 1993, *PASJ*, 45, 311