Astrophysics using Special-Purpose Computer

Junichiro Makino* and Makoto Taiji
*Department of Graphics and Information Science,
Department of Earth Science and Astronomy,
College of Arts and Sciences, University of Tokyo,
Tokyo 153, Japan
Email: makino@chianti.c.u-tokyo.ac.jp

March 2, 1996

Abstract

In theoretical astrophysics, computer simulations have been extensively used to perform experiments on astronomical objects such as stars, galaxies, and even the whole universe. Until recently, such numerical experiments had been performed on general-purpose computers. However, recent advance in the silicon VLSI technology make it practical for astrophysists to build custom computing devices to solve their own problems. As an example, we give a brief overview of our GRAPE (GRAvity PipE) project. The aim of the GRAPE project is to develop hardwares that can run astrophysical N-body simulations much faster than high-performance general-purpose computers. Our newest hardware, the GRAPE-4 system, achieved the theoretical peak speed of 1.08 Teraflops, and an N-body simulation on one-quarter of GRAPE-4 received 1995 Gordon-Bell Prize. It is world's fastest computer for scientific calculation. We describe the results of simulations on GRAPE-4 and discuss the future of GRAPE project.

Table 1: Number of particles in N-body simulations

Objects	What we can now	What is necessary
Globular clusters	10^4	10^6
Galaxies	10^5	10^7
Large-scale structures	10^7	$> 10^9$

1 Astrophysical N-body simulation

Many astronomical objects are well approximated as the system of self-gravitating particles. The solar system, open clusters, globular clusters, galaxies, clusters of galaxies are all expressed as the system of N particles interacting only through gravity.

N-body simulations have played important roles in the study of planetary formation, evolution of globular clusters, galaxies, clusters of galaxies and large scale structures of the Universe. However, even with present fastest and largest computers, what we can do is still very limited. For example, in order to understand what is going on in the center of globular clusters, in which close interaction between stars form exotic objects like X-ray binaries and millisecond pulsars, we must follow the evolution of the whole cluster, composed of one million stars, for the age of the universe (10 billion years). At present, the largest number of particles that have been used for simulation of globular cluster is 10,000, which took two months of CPU time (3 years of real time) on a Cray YMP supercomputer. Though numerical experiments with small number of particles give many important informations, large experiments with the number of particle close to that of the real system is necessary for the understanding of these systems [7]

N-body simulations are actively used to study the behavior of interacting galaxies [3] and formation of galaxies. In these simulations, we can now handle 10^5 or more particles. This number is much larger than what we can use for globular clusters (10^4), because the orbital timescale of stars is much longer for galaxies. When we simulate the formation of the large scale structure of the universe, we can use even larger number of particles (10^7 or more), since the timescale of the formation of the large scale structures is the age of the universe itself. As is shown in table 1, however, for all these systems the number of particles we can handle now is smaller than what we need by several orders of magnitude.

2 Special-purpose computer

There are several ways to perform simulations larger than what is possible on present-day computers. The simplest one is just to wait. In the last 50 years, computers have achieved truly incredible improvement in speed. The speed of the fastest computers has increased by a factor of 100 in 10 years. The price-performance has been improved in a similar rate.

Another approach is to improve numerical algorithm. The invention of the individual timestep algorithm[1] and $O(N \log N)$ tree algorithm[2] have reduced the calculation cost drastically and made it possible to use the number of particles much larger than what is possible with the simple direct algorithm.

We adopted the third approach, namely to develop the computer to solve our problem. One might think it is not the most efficient way of using their time for astrophysists to develop their own computer. There are so many computer experts, both in academic and in industry, who are working hard to develop fast computers. Well, figure 1 proves the opposite. We started our project to develop computers to solve N-body problem in 1989. In that year we completed the first working machine, GRAPE-1, with the peak speed of 240 Mflops. The name GRAPE stands for GRAvity PipE. It calculates the gravitational force between particles by specialized pipeline processors. In 1991 we completed GRAPE-3[16] with the peak speed of 15 Gflops, and in 1995 we completed GRAPE-4, the first computer to achieve the peak speed of 1 Tflops. A Simulation of binary black holes in galactic core on this system achieved the sustained speed of 529 Gflops. The 1995 Gordon Bell prize was awarded to our early calculation of a similar problem which was performed on 1/6 of the present Teraflops system and achieved the speed of 112 Gflops. In addition, copies of GRAPE-3 were now being used in more than 20 research institutes all over the world, providing the total of more than 150 Gflops of computing power.

In the following, we briefly describe the architecture of GRAPE systems and show some of scientific results recently obtained on GRAPE-4. Finally, we will discuss the future of the scientific computing.

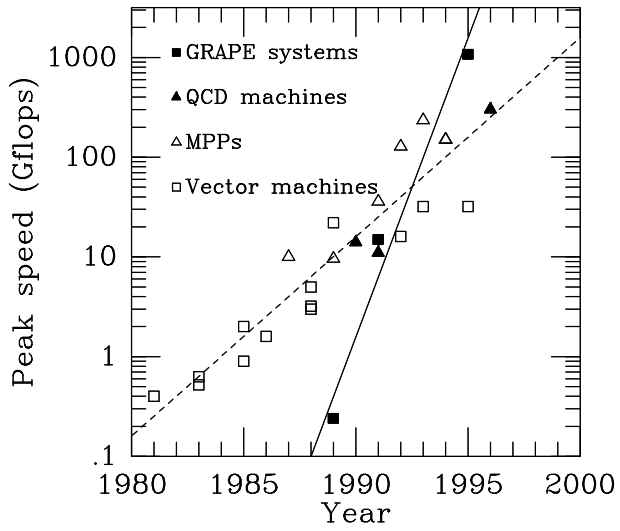

Figure 1: Peak speed of high-performance computers plotted against the date of first delivery. Filled squares, filled triangles, open triangles and open squares represent GRAPE systems, Machines specialized to QCD calculation, commercial MPPs and commercial vector (-parallel) machines. Dashed line indicates a factor of 100 improvement in 10 years, which is a good approximation for machines other than GRAPEs. Solid line represents a factor of 10^6 improvement in 10 years.

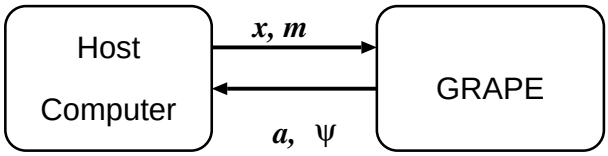

Figure 2: Basic structure of a GRAPE system

3 GRAPE system

In N-body simulations, practically all CPU time is spent to calculate the gravitational forces on particles. In the simplest algorithm, the force on a particle is calculated by summing up the forces from all other particles in the system. Thus the cost of force calculation is $O(N^2)$, while the cost of the rest is $O(N)$.

Figure 2 shows the basic structure of our GRAPE systems[18]. The special-purpose hardware calculates the interaction between particles. The host computer (a workstation) performs all other calculations and I/O. We call the special-purpose hardware GRAPE (GRAvity PipE), since its main computing unit is the hardwired pipeline to calculate the gravitational force between particles.

The direct force calculation is well suited for acceleration by specialized hardware, because of its simplicity. We calculate the difference in the positions of two particles $\Delta \mathbf{r}_{ij}$, and calculate the Newtonian gravitational force, $-Gm_j \Delta \mathbf{r}_{ij}/|\Delta r_{ij}|^3$, where G is the gravitational constant and m_j is the mass of particle j. With present VLSI technology, a fully hardwired pipeline to calculate the gravitational force can be implemented on a custom LSI chip. It is also possible to make a massively parallel machine with thousands of custom chips. The GRAPE-4 system integrates 1692 pipeline processor chips. Each chip integrates some 20 floating point units and operates at 32 MHz. Thus the peak speed of a chip is 640 Mflops and that of the total system is 1.08 Tflops.

In the simplest case, the host computer sends the positions and masses of all particles to GRAPE. Then GRAPE calculates the force between particles, and sends the forces on particles back to the host computer.

From the software point of view, all program runs on the host. The

special-purpose hardware is used through the calls to a few library functions, which calculate the force between particles. It is not very difficult to modify existing programs so that it uses the special-purpose hardware, since the task is the same as to utilize a software library[21], [4]. To develop the library functions themselves is not very difficult either, since the number of functions is very small[13].

To develop the hardware is not a very difficult task either. The total number of people who was involved in the design of GRAPE-4 is three and it took two years to complete the design. This investment is by no means larger than the investment on the simulation softwares.

4 GRAPE and sophisticated algorithms

The direct summation algorithm, in which we calculate the force on a particle by actually summing up the forces from all other particles, is obviously not the most efficient algorithm. For some simulations, fast algorithms such as Barnes-Hut tree is used. We decided to develop the hardware to accelerate the direct force calculation because of the following reasons.

First, many applications require numerical accuracy which is difficult to achieve by fast algorithms. Second, a specialized hardware for pairwise force calculation can be used to improve the speed of sophisticated algorithms such as Barnes-Hut tree [11] and Particle-Particle Particle-Mesh (P^3M) scheme [4]. In both schemes, the speedup by a factor larger than 10 was achieved.

We could have designed a machine specialized to sophisticated schemes such as the Barnes-Hut tree algorithm. We did not adopt this approach since our main interest lies in the problems for which sophisticated schemes are difficult to apply.

5 GRAPE-4

5.1 Architecture

The GRAPE-4 system consists of a host computer and four clusters(figures 4,3). One cluster has one host-interface board (HIB), one control board (CB), and 9 processor boards (PB). HIB and CB handles the communication between the processor boards and the host. The processor boards perform the

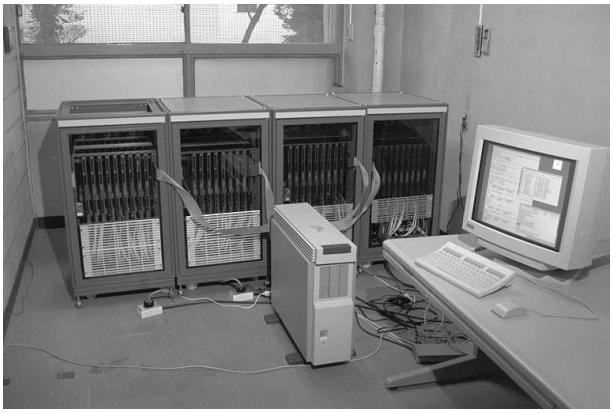

Figure 3: The GRAPE-4 system with its host computer DEC Alpha AXP 3000/900.

force calculation. Each processor board houses 48 HARP (Hermite AccceleRator Pipeline) chips, which are custom LSI chips to calculate the gravitational force and its first time derivative. The theoretical peak speed of GRAPE-4 is 1080 Gflops, for the clock frequency of 32 MHz.

5.2 Parallelism and calculation flow

The total number of pipeline chips of GRAPE-4 is 1692. Each pipeline calculates the forces on two particles, using the technique we call "virtual multiple pipeline". The simplest algorithm to use this number of chip is to let them calculate forces on 3384 different particles. With this approach, however, the degree of parallelism is too large to run the individual timestep algorithm efficiently, unless the total number of particles is unpractically large. Therefore, we used a different approach: Each of the 36 processor boards calculates the forces on same 94 particles, but from different set of particles. Thus, the total force on a particle is computed as 36 partial sums. The summation within a cluster is performed on CB and summation of contributions from different clusters is performed on the host computer.

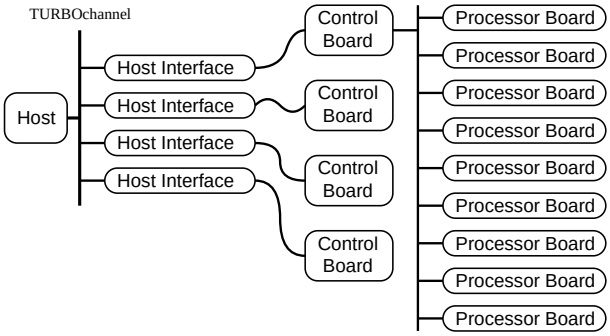

Figure 4: The GRAPE-4 system

5.3 Processor board

A processor board consists of the particle data memory, one predictor LSI chip and 47 HARP LSI chips. The particle data memory stores the data of particles which exert the force. The predictor LSI is used to calculate the position (and velocity) of particles at a specified time. This function is necessary to use the individual timestep. The HARP LSI chips calculate the gravitational accelerations and their first time derivatives for particles. One board has 48 HARP chips. 8 HARP chips are packaged in one custom MCM (multi-chip module) package(figure 5). One PB board houses 6 MCMs. We use 47 out of 48 chips for actual calculation in order to utilize MCMs with one defect chips. Each HARP chip calculates the forces on two different particles simultaneously. Thus one board calculates the forces on 94 particles in parallel.

5.4 Hardware cost

The total cost of GRAPE-4 is 150 M JYE, of which 80 M JYE is spent for the production of the hardware. Remaining 70 M JYE is spent for the designing (including the cost of design software and workstations to run the software). The cost per Gflops is 0.17 M JYE.

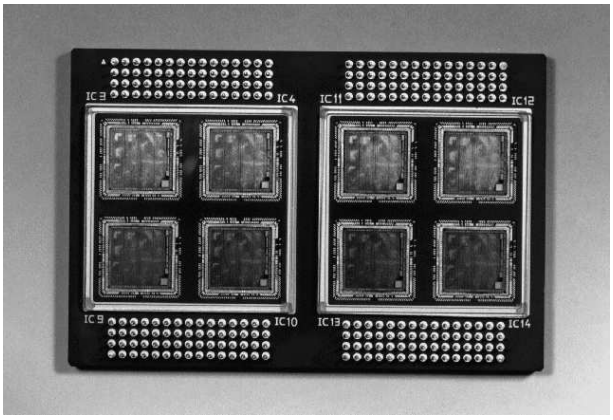

Figure 5: The MCM (multi-chip-module) package which houses eight HARP pipeline chips.

6 Black hole binary at the center of the galaxy

Recent high-resolution observation by Hubble Space Telescope (HST) have revealed unexpected fact that most of large elliptical galaxies have the "density cusp" at the center, in which the stellar density ρ is proportional to the inverse of the distance from the center r [**?**]. Standard scenarios of the formation of ellipticals cannot explain these cusps.

Even before this new finding by HST, the central region of large ellipticals had been difficult to understand. Figure 6 shows the relation between the "core radius" and the effective radius (radius which contains half of the total luminosity) of ellipticals[9]. Here, core radii are determined by ground-based observation. It is clear that the "core" radius and the effective radius shows correlation.

The standard scenario of forming large ellipticals is merging of smaller galaxies. It nicely explains many of the observed characteristics of elliptical galaxies. Early simulations of merging galaxies [5], however, have demonstrated that if progenitors have cores of similar sizes, the merger remnant also have the core of almost the same size. On the other hand, the effective radius becomes larger by approximately a factor of two. Thus, the merger scenario cannot reproduce the relation between the core size and the effective

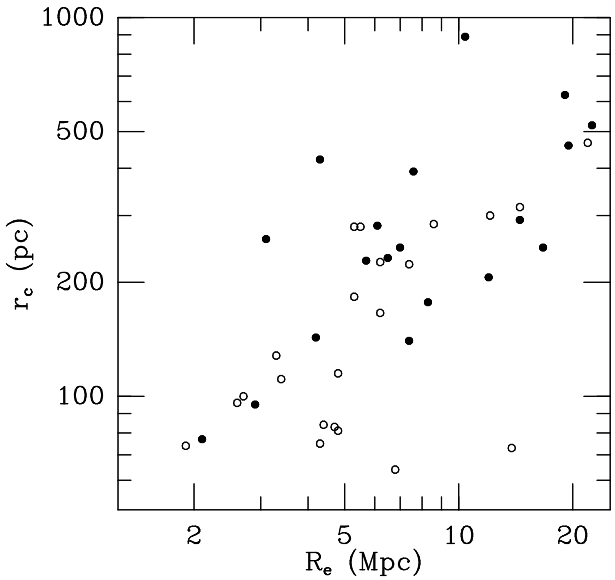

Figure 6: Relation between the core radius (vertical axis) and the effective radius (horizontal axis) of observed elliptical galaxies. Filled circles indicates the galaxies of the resolution classes I and II in Lauer (1985). Open circles denote those of class III.

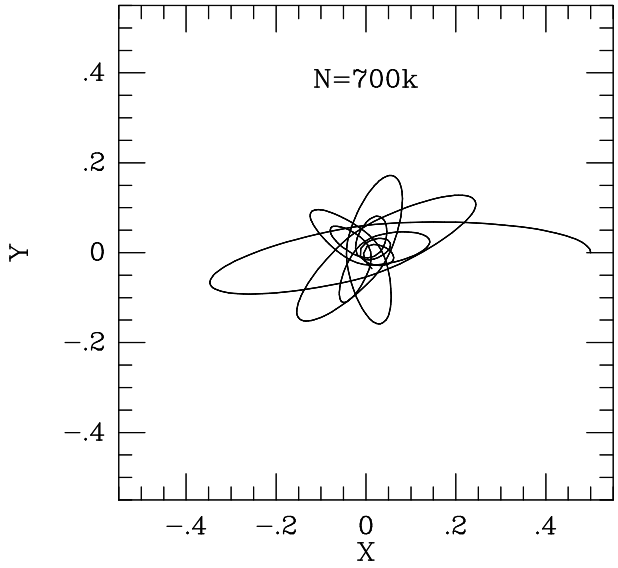

Figure 7: The trajectory of one binary in x-y plane. The parent galaxy is a spherical Plummer model with the half-mass radius of 0.77.

radius.

Many ellipticals are believed to have central massive black holes with the mass comparable to the estimated core mass itself [10]. However, in simulations by Farouki et al. [5], the effect of black holes was ignored.

What will happen if we have two black holes in a galaxy? Figure 7 shows the result of calculation that we performed on GRAPE-4. This calculation used $700,000 + 2$ particles and took about one day on full 4-cluster configuration of GRAPE-4, running at the sustained speed of 529 Gflops. In this calculation, we put two black holes in an initially spherical galaxy. These black holes sink into the center of the galaxy. This is because of the effect called "dynamical friction". If the velocity of the black hole is smaller than the typical orbital velocity of stars around the black hole, the frictional force is proportional to the velocity of the black hole multiplied by its mass. In typical ellipticals, the timescale of the dynamical friction is similar to the orbital timescale. Thus black holes sink toward the center rather quickly.

Through dynamical friction the stars in the galaxy receive kinetic energy

from black holes. Thus, when black holes reach to the center of the galaxy, the distribution of the stars will be significantly affected. In order to investigate the structure of the central region, we did the simulation similar to that of Farouki et al. [5], but with and without central black holes [14]. Figure 8 shows the density profiles of the mergers. The radius and density are scaled so that the half-mass radius is the same for all mergers. For simulation without BH, figure 8 shows that the core size relative to the half mass radius becomes smaller, while it is almost constant for merging with central BHs. These are calculations using about 32k particles. We are currently running simulations with 256k particles to obtain more accurate central structure.

7 Other results

GRAPE-4 is now being used extensively on several projects. Here we just briefly overview them.

One of the main scientific targets of GRAPE-4 system is to simulate the secular evolution of globular clusters with more than 20,000 stars. Simplified models predicted that central density would show nonlinear oscillation with a large amplitude (more than three orders of magnitudes) [20]. For this problem, the largest number of particles ever tried on general purpose computer is 10,000 [17]. Here, the number of particles is small since calculation cost is known to be $O(N^{3.3})$ [12]. With GRAPE-4, we calculated the evolution of a 32768-body system and obtained clear proof that this oscillation actually occurs [14].

This calculation, which took three months on one quarter of GRAPE-4 system running at the speed of 50 Gflops, is probably the largest single scientific calculation ever done.

Yoko Funato has done systematic survey of the effect of the encounters between two galaxies[6]. Close encounters between galaxies are rather frequent events in clusters of galaxies, since the average distance between galaxies is not much larger than the size of a typical galaxy. However, how the global quantities (for example the mass and radius) and the structure of the galaxies have not been well understood, since we have to perform many experiments to determine the dependence of the result on the collision parameters such as the relative velocity and impact parameters. Funato did several hundreds of scattering experiments to determine the effect.

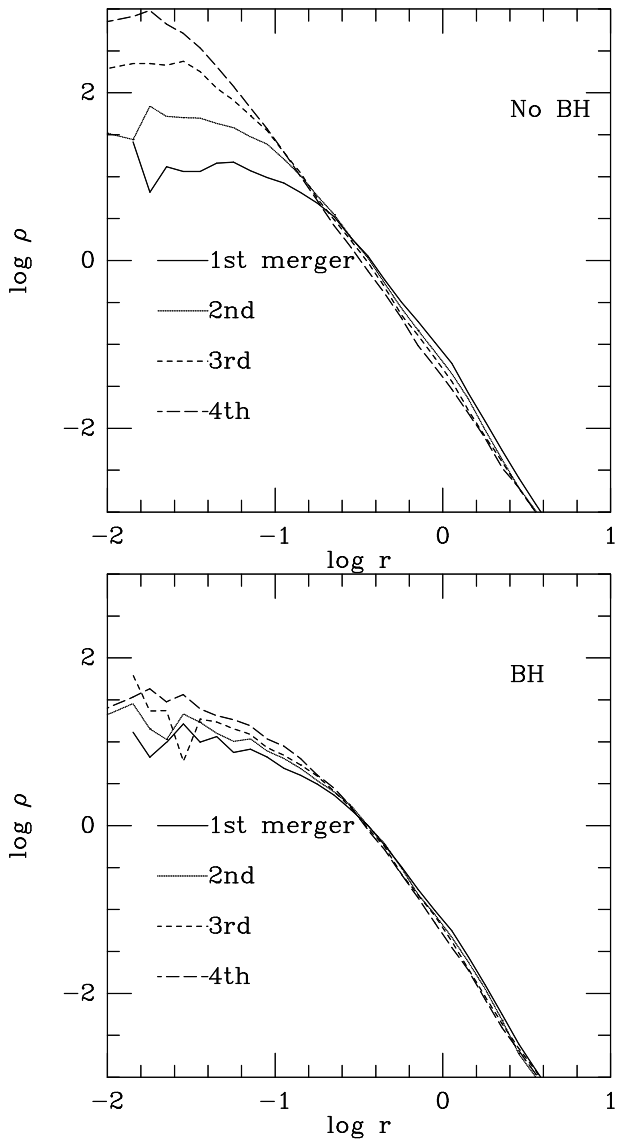

Figure 8: The density profiles of the mergers (a) without and (b) with central black holes. Radius and density are scaled so that the half-mass radius is the same for all merger remnants.

Eiichiro Kokubo and Shigeru Ida have been studying the formation process of planets using several generations of GRAPE hardwares. In standard scenario of planetary formation, planets are formed by coagulation of smaller objects called planetesimals. Thus, in order to understand the formation process, we have to perform the simulation of the system of many planetesimals. Recently, they have done full three-dimensional simulation of the system of 8,000 planetesimals. Theirs is the first simulation that handles the spatial inhomogeneity correctly, and thus the first step to explain the location of planets (Bode's law).

Toshiyuki Fukushige is studying the formation process of the dark matter halos of galaxies. This is an important problem which might be the key to understand the discrepancy between theoretical predictions and observations of large scale structures of the universe. There are several studies on this problem [15], but the accuracy of their results is severely limited by the number of particles employed ($< 3 \times 10^4$). With GRAPE-4 system it is possible to handle 5×10^5 or more particles, which will provide necessary mass resolution to model the initial power spectrum of the density fluctuation of the early universe.

8 The Future

In the past, the evolution of the computational physics, or the computational science in general, has been in large part driven by the evolution of the computer hardware. Of course, the development of sophisticated algorithms also played significant role, but without the increase of the speed of computers by a factor of 10^{10} in just a half century, the computational science would not have existed.

However, whether or not such an advance in the computer technology will continue or not is rather unclear at present, since almost all manufacturers of supercomputers and massively parallel computers ceased to exist. The silicon VLSI technology will continue to evolve, but there might not be any company to make use of it to develop fast and large computers except for big Japanese companies.

This uncertainty comes partly from political reasons, but mainly from limitations in the computer architecture.

The advance in the semiconductor technology has two consequences. The

first one is the improvement in the switching speed, which is directly reflected to the speed of the computer. The second one is the increase in the available number of transistors, which is more difficult to utilize.

The evolution of the floating-point performance of microprocessors illustrates the problem. In early 1980s, first single-chip floating-point processor chips like Intel 8087 were developed. These chips have the speed less than 0.1 Mflops. In 1989, Intel started to ship i860, which has the peak speed of 80 Mflops. Thus, in less than 10 years the improvement of a factor close to 1,000 was achieved. In 1990s, however, the improvement slowed down. The fastest microprocessors currently in production offer the speed of 700-800 Mflops, which is only a factor of 10 improvement in 7 years.

In 1980s, the increase in the transistor count contributed to the improvement in the speed of floating point unit, through the reduction of CPI (cycles per instruction). A pipelined floating point multiplier for 64-bit numbers requires about 10^5 transistors. If this number of transistors cannot fit in a chip, we have to use multiplier which needs several clock cycles to produce one result. In other words, as far as the available number of transistors for multiplier is less than 10^5, the increase in the number of transistor can be used to improve the speed of a single multiplier. This is the reason why the improvement in the speed was so rapid in 1980s. Intel i860 is the first single-chip microprocessor to integrate fully-pipelined floating-point unit. The clock speed of i860 was 40 MHz, which is only 10 times faster than the clock speed of 8087. Even so, i860 is nearly 1000 times faster than 8087.

To utilize multiple floating point units within one microprocessor is not easy. As a result, Microprocessors of 1996 still have only one or two multipliers, instead of 10 or 20 which is theoretically possible.

It should be noted that this difficulty in using multiple functional units in microprocessors is exactly the same difficulty as mainframe and supercomputer designers saw 25 years ago. CDC-7600 is one of the first machines to have fully pipelined floating-point unit. The evolution of microprocessor is an accelerated replay of the evolution of the computer architecture itself.

There are two reasons why it is difficult to use multiple floating point units. One is a software problem. It is not easy to find sufficient degree of parallelism to keep multiple functional units busy. This is especially true if the processor is designed to keep backward compatibility with its predecessors. The other is a hardware problem. It is very difficult to increase the bandwidth between the memory and the processor, since the speed of I/O

and the number of I/O pins are more difficult to increase than the speed of the signal the number of transistors in a chip. Thus, the data transfer bandwidth of present-day microprocessors is only 2-3 times faster than that of i860.

Special-purpose machines like GRAPE have none of these problems. In GRAPE architecture, we can use about 60 floating point units to construct one pipeline for force calculation. If the available number of transistors is larger than that is needed for a single pipeline, we can put multiple pipelines into one chip. There is no software problem since all the floating point units keep doing the same operations. There is no hardware problem (bandwidth problem) either, since multiple pipelines can share the same input data just in the same way as the multiple pipelines on GRAPE-4 processor boards share the same input data.

Special-purpose systems like GRAPE can continue to improve the performance in the rate similar to that achieved by microprocessors in 1980s (a factor of 1000 in 10 years). On the other hand, microprocessor-based systems have already slowed down the evolution by 1990. Thus, the relative advantage of special-purpose systems over the general-purpose computers will continue to increase, unless the design principle of general-purpose computers will change drastically.

9 Towards Petaflops

A factor of 1000 in 10 years means we will be able to build a Petaflops system by 2005, provided that we can get the funding of similar amount as we got for GRAPE-4. We are now trying to achieve the Petaflops speed before the end of this century. The Petaflops GRAPE will consist of about 20,000 pipeline processors, each delivering the speed of 50 Gflops. This speed will be realized by packing 300 floating arithmetic units running at 160 MHz clock (mainly single precision) into one chip. This goal will not be very hard with the 0.25 μm technology that will be available by 1998.

The total development cost will be about 10 M dollars. We are currently seeking for funding, both within and outside Japan.

General-purpose computers will achieve Petaflops by around 2010 Thus, Petaflops GRAPE will be delivered 10 years in advance. For some problems, this difference in time will be even larger, since many algorithms are difficult

to run efficiently on massively parallel processors. If a Petaflops general-purpose computer will be delivered by 2010, it will certainly be an MPP with more than 100,000 processors. Such a machine will be useful only for rather limited class of problems, and the astrophysical N-body problem might not be one of them.

References

[1] S. J. Aarseth, *Monthly Notices Roy. Astron. Soc.*, Vol. 126, p. 223, 1963.

[2] J. E. Barnes and P. Hut, *Nature* Vol. 324, p. 446, 1986.

[3] J. E. Barnes and L. Hernquist, *Annual Review of Astronomy and Astrophysics*, Vol. X, p. X, X.

[4] P. Brieu, F. Summers, and Ostriker, J. P., *Astrophys. J.*, Vol. 453, p. 566, 1995.

[5] R. T. Farouki, S. L. Shapiro, and M. J. Duncan, *Astrophys. J.*, Vol. 265, 597, 1983.

[6] Y. Funato, PhD Thesis, University of Tokyo, 1996.

[7] D. C. Heggie, in in *Dynamical Evolution of Star Clusters — Confrontation of Theory and Observation, Proceedings IAU Symposium 174*, eds P. Hut and J. Makino, Kluwer Academic Publishers, Dordrecht, 1996 (in press).

[8] T. R. Lauer, *Astrophys. J.*, Vol. 292, 104, 1985.

[9] T. R. Lauer, E. A. Ajhar, Y. I. Byun, A. Dressler, S. M. Faber, C. Grillmair, J. Kormendy, D. Richstone, and S. Tremaine, *Astron. J.* Vol. 110, p. 2622, 1995.

[10] J. Kormendy and G. Djorgovski, Annual Review of Astronomy and Astrophysics, Vo. 27, 235, 1988.

[11] J. Makino, *Publ. Astron. Soc. Japan*, Vol. 43, p. 621, 1991.

[12] J. Makino and P. Hut, *Astrophys. J. Suppl.*, Vol. 68, p. 833, 1988.

[13] J. Makino and Y. Funato, *Publ. Astron. Soc. J.*, Vol. 45, p. 279, 1993.

[14] J. Makino and T. Ebisuzaki, *Astrophys. J.* (in press).

[15] J. F. Navarro, C. S. Frenk and S. D. M. White, submitted to *Astrophys. J.*.

[16] S. K. Okumura, J. Makino, T. Ebisuzaki, T. Ito, T. Fukushige, D. Sugimoto, E. Hashimoto, K. Tomida and N. Miyakawa, *Publ. Astron. Soc. J.*, Vol. 45, p. 329, 1993.

[17] R. Spurzem and S. Aarseth, submitted to *Mon. Not. Roy. Astron. Soc.*.

[18] D. Sugimoto, Y. Chikada, J. Makino, T. Ito, T. Ebisuzaki and M. Umemura, *Nature*, Vol. 345, p. 33, 1990.

[19] D. Sugimoto, Y. Chikada, J. Makino, T. Ito, T. Ebisuzaki and M. Umemura, *Nature*, Vol. 345, p. 33, 1990.

[20] D. Sugimoto and E. Bettwieser, *Mon. Not. Roy. Astron. Soc.*, Vol. 204, 19p, 1983.

[21] M. Umemura, T. Fukushige, J. Makino, T. Ebisuzaki, D. Sugimoto, E. L. Turner, and A. Loeb, *Publ. Astron. Soc. J.*, Vol. 45, p. 311, 1993.