GRAPE-4 and Beyond

Junichiro Makino *

Abstract

We briefly describe the past, the present and the future of the GRAPE project. GRAPE is specialized computer to perform N-body simulation. We have completed the GRAPE-4 system in summer 1995. It is the first computer to achieve the peak speed exceeding 1 Teraflops. We overview some of scientific works being done on GRAPE-4. It is currently used to study various topics such as the pre- and post-collapse evolution of globular clusters, evolution of binary black holes in galactic nuclei, clusters of galaxies, formation of galactic halos, and systems of planetesimals. Finally, We discuss the follow-up project, the Petaflops GRAPE. To develop a Petaflops GRAPE system before the end of century is technologically feasible. It will enable us to tackle the problems that cannot be solved on general-purpose machines before 2010.

1. The past

We started to develop GRAPE machines in late 1988, following the proposal which Chikada presented in 1988 Summer school of young astronomers. In his proposal, Chikada urged theoretical astrophysists to design and build their own digital hardware to do numerical computation. At that time, we had been using vector machines and various kind of parallel machines for several years, and were frustrated by the limitation of the machines available then and the direction of the evolution of the hardware then predicted.

So we started to develop hardware. We do not give the description of the basic idea, which can be found elsewhere[1]. To our surprise, so far we have been quite successful. Figure 1 shows the evolution of GRAPE systems compared to that of general-purpose parallel/vector computers and machines dedicated to QCD calculation. In 1995 summer, we completed GRAPE-4 system[8] with the peak speed of 1.08 Tflops. The simulation of blackhole binary in galactic core won the 1995 Gordon Bell prize, an award given to the calculation which achieved the fastest measured performance.

2. The present

GRAPE-4 is currently being used on various problems. Here we give a few examples.

One of the goals of GRAPE-4 project was to develop a machine fast enough to determine if the gravothermal oscillation[7] occurs in N-body systems. For this purpose, we need to follow the evolution of a system with more than 20K particles for several thousand crossing times. The

*Department of Graphics and Information Science, College of Arts and Science, University of Tokyo, 3-8-1 Komaba, Meguro-ku, Tokyo 153, Japan

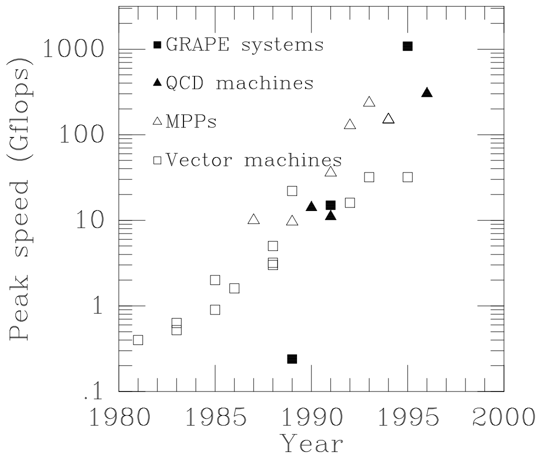

Figure 1. The peak speed of GRAPE systems (filled squares), QCD machines (filled triangles), MPPs (open triangles) and vector machines (open squares) are plotted against the year of first delivery.

largest calculation ever tried on conventional vector machine used 10K particles. On a one-quarter of hardware, we have run 32K-particle simulation and obtained clear evidence that the gravothermal oscillation takes place in N-body systems[4].

We have been studying the evolution of blackhole pair in galactic core. It may play an important role to determine the structure of the central core of large ellipticals, as recently revealed by HST observation[3]. Our 32k-particle simulation[5] has shown that the presence of blackhole pair naturally reproduce the observed structure of central cores of large ellipticals. We are currently investigating the structure of the core with higher resolutions (256K and 512K particles)

Funato[2] studied the effect of close encounters between galaxies in clusters of galaxies. She modeled galaxies by 4K particle. She found that the evolutionary track of galaxies in the plane of luminosity-velocity dispersion depends on the power index of halo density, and that the Faber-Jackson relation is realized for the "natural" profile of $\rho \propto r^{-4}$.

Kokubo and Ida have been studying the growth of protoplanets by means of direct N-body simulation, using GRAPE-4 and its prototype hardware, the HARP-2.

3. The Future

We have started to develop the next generation machine. Here we briefly describe the schedule, the performance, and the cost of the next-generation machine.

At present, our plan is to develop a petaflops GRAPE by the year 2000. This plan might sound too ambitious. However, it is perfectly feasible with the technology available by 1998. GRAPE-4 is based on 1μm technology, which was available in 1990. We started the development of GRAPE-4 in 1992 and made a rather conservative choice to use the technology already available at that time. As a result, the pipeline processor chip of GRAPE-4 consists of only 100k gates and operates at rather slow clock of 32 MHz.

In 1995, a few Japanese companies started to produce 0.35μm devices. With this technology, we can pack 1M gates operating at 100 MHz clock in a single chip. Such a chip will achieve 30 times more speed than present HARP chip used on GRAPE-4.

In 1998, most LSI companies will be producing 0.25μm devices and some will proceed to 0.18μm. Such technology will enable us to pack 3M or more gates in a single chip which will operate at the clock frequency of over 200 MHz. In other words, GRAPE-type chip of 1998 will be about 200 times faster than the present HARP chip. Thus, in order to achieve one Petaflops, we need only five times more pipeline chips than is used on GRAPE-4.

The power consumption of Petaflops system will not be excessive. GRAPE-4 dissipates about 10KW. The power consumption of Petaflops system will be about 25 times as that of GRAPE-4, since the speed is 1,000 times faster and power consumption per unit speed will be about 1/40. This reduction in power consumption per speed comes from the reduction in the feature size and the power supply voltage. The power consumption of 250 KW is large, but is still similar to that of a typical vector supercomputer.

The development cost of a Petaflops machine will be about one billion yens (10M dollars), which is five times more than that of GRAPE-2 (185 M yens). We are currently seeking for funding possibilities both within and outside Japan.

The speed of general-purpose computer is improved by only a factor of 10 in five years. We believe we can achieve a factor of 1000 in the same five years because of the following three reasons. First, a GRAPE system can utilize all the available transistors on an LSI for floating point units, while on general-purpose computers increasing fraction of total transistors is used for cache memories and control logics. Thus, we can achieve a factor of 30 in five years instead of a factor of 10. Second, the time difference in technology will be 8 years if we will adopt the technology of 1998, since GRAPE-4 is based on technology of 1990. Thus, we can achieve a factor of 200 improvement in price-performance. The last factor of five will (hopefully) come from larger budget.

The P.E.T.A. group predicted the delivery of the first petaflops machine to be around 2013[6]. Once the budget is approved, petaflops GRAPE will definitely be the first computer to achieve the petaflops.

References

[1] Ebisuzaki, T., Makino, J., Fukushige, T., Sugimoto, D., Ito, T. & Okumura, S. K. 1993, PASJ, 45, 269.

[2] Funato, Y. 1996, Ph. D. Thesis, University of Tokyo, Japan.

[3] Lauer, T. R., Ajhar, E. A., Byun, Y. I., Dressler, A. Faber, S. M., Grillmair, C., Kormendy, J. Richstone, D. & Tremaine, S. 1995, AJ, 110, 2622.

[4] Makino, J. 1996, in Dynamical Evolution of Globular Clusters, Proceedings of IAU Sympsium 174, eds P. Hut and J. Makino, Kluwer (in press).

[5] Makino, J. & Ebisuzaki, T., 1996, ApJ, 465 (in press).

[6] Sterling, T., Messina, P. & Smith, P. H. 1995, Enabling Technologies for Petaflops Computing, MIT Press, Cambridge.

[7] Sugimoto, D. & Bettweiser, E., 1983, MNRAS, 204, 19p.

[8] Taiji, M., Makino, J., Fukushige, T., Ebisuzaki, T. & Sugimoto, D. 1996, in Dynamical Evolution of Globular Clusters, Proceedings of IAU Sympsium 174, eds P. Hut and J. Makino, Kluwer (in press).