A 992 Gflops simulation of black holes in a galactic center on GRAPE-6

Junichiro Makino*, Toshiyuki Fukushige and Masaki Koga
*Department of Astronomy, School of Science, University of Tokyo,
Tokyo 113-0033, Japan
Email: makino@astron.s.u-tokyo.ac.jp
Department of General System Studies,
College of Arts and Sciences, University of Tokyo,
Tokyo 153-8902, Japan

Abstract

As an entry for the 2000 Gordon Bell performance prize, we report the performance achieved on a prototype GRAPE-6 system. GRAPE-6 is a special-purpose computer for astrophysical N-body calculations. The present configuration has 64 custom pipeline processors, each containing six pipeline processors for the calculation of gravitational interactions between particles. Its theoretical peak performance is 1.926 Tflops. The complete GRAPE-6 system will consist of 3072 pipeline chips and will achieve a peak speed of 100 Tflops. The actual performance obtained on the present 64-chip system was 992 Gflops, for a simulation of massive black holes embedded in the core of a galaxy with 786,432 stars.

1 Introduction

The structure and dynamics of the central region of galaxies is an important field of research in astrophysics. Of particular interest is the existence of massive black holes. Observations using the Hubble Space Telescope (*e.g.*, [vv98]) and large interferometer arrays (*e.g.*, [MMH$^+$95]) have shown that many galaxies host central massive objects with a mass million or even billion times larger than that of the mass of our sun.

The most direct observational evidence for the existence of a massive black hole is for the presence of a black hole in the very core of our own galaxy [EOG99]. Here, thanks to our close proximity, observers can now see individual stars with a distance of less than 0.01 parsecs from the central massive object and measure their radial (line of sight) velocities as well as tangential velocities (through proper motion measurements). All evidence is quite consistent with the existence of a massive black hole with a mass of $2.7 \times 10^6 M_\odot$ (here $M_\odot$ is one solar mass) at the center of the galaxy.

The evidence for the existence of black holes in distant galaxies is less direct. Nonetheless, there is now little doubt that most galaxies host massive central objects, which are most likely massive black holes. Among these massive black holes, those in large elliptical galaxies are known to be the most massive. The mass of black holes in normal spiral galaxies is typically several million solar masses. The mass of black holes in elliptical galaxies is much larger, typically more than a hundred million solar masses.

How are such massive black holes formed? How is their formation process related to the formation process of the galaxy itself? Why do elliptical galaxies have more massive black holes than spiral galaxies? In order to answer these questions, we need to interpret observations with the help of numerical experiments (simulations).

The standard theory of the formation of elliptical galaxies is that they are formed through merging of smaller galaxies, either spiral galaxies or elliptical galaxies [TT72]. Numerical simulations of galaxy merging have demonstrated that this theory explains many observed features of elliptical galaxies very well.

When two spiral galaxies merge, the gas in the disk of a galaxy is strongly perturbed by the gravitational force from the other galaxy, and tends to concentrate into the central region of the galaxy well before two galaxies actually merge. Such a high-density gas results in very high star forming activity, which is observed in many interacting galaxies. Also, some of the gas would eventually be absorbed by the central black hole.

Thus, when two galaxies finally merge, the black holes have become much more massive than their original mass, and most of the gas in the galaxies is lost. These black holes will settle in the center of the merger galaxy, through the dynamical friction from other stars in the galaxy. In other words, the galaxy that is the product of the merger of two galaxies is likely to have two black holes in the center. What will happen to these two black holes?

These black holes form a bounded pair (binary), and the binding energy of the binary increases over time through dynamical friction. If the dynamical friction remains effective until the separation of the two black holes in the binary becomes small enough for the gravitational radiation to shrink the separation further, the two black holes will merge. Such mergings of two massive black holes would emit strong gravitational wave which will be observable through space-based gravitational wave detectors such as LISA [JCR94].

However, N-body simulations of the evolution of binary black holes in the center of a galaxy [ME96, Qui96, Mak97, QH97] have suggested a different picture. Before the separation of the black holes becomes small enough, the binary black hole expels almost all the stars in its vicinity, resulting in a region devoid of stars. As a result, dynamical friction from field stars becomes inefficient.

N-body simulations of binary black holes in the center of the galaxy pose numerically very tough problems. There are several reasons why these are tough. In order to obtain a reliable result, we need to use as many particles (stars) as possible, to reduce the numerical relaxation effects. Here, numerical relaxation denotes the effect that stars change their velocity randomly through encounters with other stars. The strength of this effect is proportional to $1/N$, where N is the number of stars in the system. In the limit of $N \to \infty$, this effect becomes zero. In real galaxies, this effect exists, but is very small. In numerical simulations with, say, 10^5 stars, this effect is several orders of magnitude larger than that in real galaxies.

Another reason is that the simulation should cover a very wide range in timescales. Typical stars in a galaxy have orbital timescales of 10^8 years, and therefore the timestep necessary to integrate the orbit of these stars accurately is around 10^5 years. The orbital period of the black hole binary can be anywhere between 10^6 years and 10 minutes. Of course, we do not need to follow a binary with an orbital period of 10 minutes for a long time, since the black holes would merge in a few orbital periods anyway. However, in order to see if gravitational radiation becomes effective or not, we need to follow binaries with an orbital period of about one year. Thus, it is crucial to assign different timesteps to different stars (using individual timestep algorithms, [Aar63]).

This requirement to handle a wide range of the timescales leads to the additional requirement of very high accuracy, simply because black holes and stars close to black holes are integrated for a very large number of orbits. The gravitational forces on particles should be calculated to high accuracy. The orbits should also be integrated with high accuracy.

These two requirements of a wide range in timescales and high accuracy makes the numerical simulation very difficult. In particular, no one has succeeded to use the tree algorithm for this kind of problem. The tree algorithm is now widely used for simulations of the formation and evolution of galaxy clusters, galaxy groups and individual galaxies. However, in these calculations the central regions of galaxies are not correctly modeled.

It is not impossible to combine the individual timestep approach with a tree algorithm [MA93], and achieve high accuracy at the same time. However, in order to obtain astrophysically useful results from such a code, parallelization is necessary in order to obtain high calculation speeds, and that poses a difficult problem.

In fact, even without tree algorithm, the implementation of the individual timestep algorithm on massively parallel computers turned out to be indeed a difficult problem. The reason is that we need a very low-latency network which can perform broadcast or global reduction, or preferably both. The fraction of the stars with very small timestep is small, but since they have such small timesteps, a large fraction of the total calculation time is spent to obtain gravitational forces on them. So we have to efficiently parallelize the calculation of the force on a small number of particles. One can get good performance on vector processors, shared-memory multiprocessors, and small-scale distributed-memory parallel computers with a very low-latency network, but with present large MPPs, it's very difficult to achieve good performance.

Black holes in the center of a galaxy form just one example of tough problems in astrophysics. Many astrophysical systems share the same characteristic of exhibiting a wide range in timescales. This problem arises directly from the fact that the gravitational force is an attractive force with no characteristic scale length.

In order to accelerate N-body simulations with individual timestep algorithms, we have developed a series of special-purpose hardwares for the force calculation [SCM+90, MT98]. The direct force calculation is well suited for acceleration by specialized hardware, because of its simplicity. The use of individual timesteps adds extra complexity to the hardware and software, but still we can achieve pretty high performance. In the case of special-purpose computers, it is not too difficult to implement a fast network for broadcast and

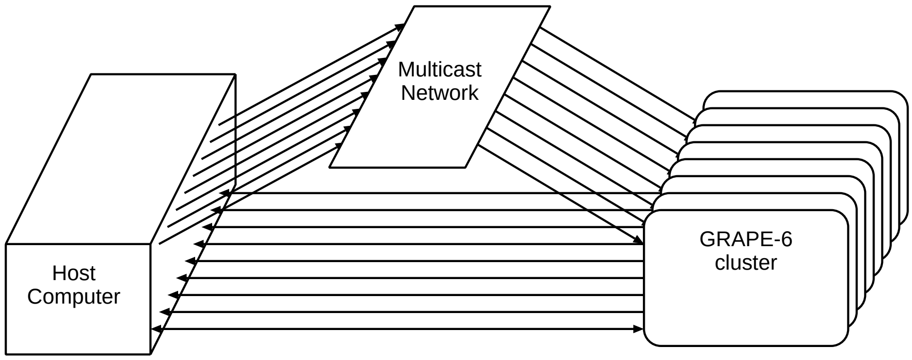

Figure 1: Overall structure of GRAPE-6 system. Eight clusters are shown.

reduction. So the latency problem of general-purpose massively parallel computers is solved by hardware.

It should be noted that the use of GRAPE hardware is not limited to direct summation. GRAPE can achieve pretty high performance for the tree algorithm, as demonstrated by our entry for the Gordon Bell prize last year.

In 1995, we completed GRAPE-4[MTES97] with a peak speed of 1.08 Tflops. The direct N-body simulations on this system were awarded the Gordon Bell prize in 1995 (112 Gflops) and 1996 (333Gflops). In 1997, we started the development of the GRAPE-6 system. The full GRAPE-6 system will be completed by year 2001, and will achieve a peak speed of 100 Tflops. In this paper, we report the performance of 1/48 of the total system, which is currently up and running as a testbed. In section 2, we describe the GRAPE-6 system and its present status. In section 3, we report the simulation performed on GRAPE-6.

2 The GRAPE-6 system

2.1 Architecture

The GRAPE-6 system will consist of a host computer and 12 clusters. Each cluster has two I/O ports, one for local communication and the other for multicast through clusters. The multicast network is necessary to implement the individual timestep algorithm efficiently, but in this paper we do not discuss the details, and concentrate on a single cluster. Figure 1 shows the overall structure of the GRAPE-6 system.

One cluster has one host-interface board (HIB), five network boards (NB), and 16 processor boards (PB). Each NB has one uplink and four downlinks. Thus, 16 PBs are connected to the host through two-level tree network of NBs (see figure 2). HIB and NB handles the communication between PBs and the host.

The PBs perform the force calculation. Each PB houses 16 GRAPE-6 processor chips, which are custom LSI chips to calculate the gravitational force and its first time derivative. Figure 3 shows the photograph of a PB. Two processor chips and four memory chips are

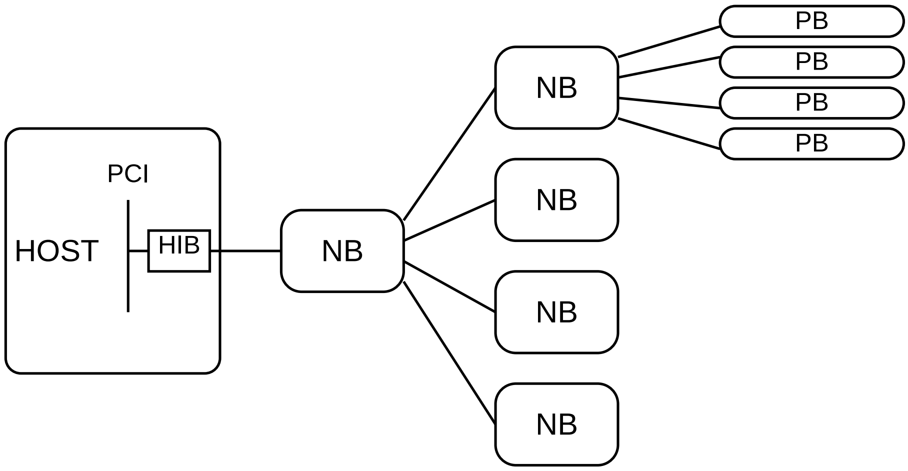

Figure 2: A GRAPE-6 cluster. First four PBs are shown. Other 12 PBs and interface to the multicast network are omitted.

mounted on a daughter card, and eight daughter cards are mounted on a processor board. Three large chips on board are the FPGA chips for the reduction network.

A single GRAPE-6 processor chip integrates six pipeline processors for the force calculation, one pipeline processor to handle the prediction, network interface and memory interface (see figure 4). One force pipeline can evaluate one particle-particle interaction per cycle. With the present pipeline clock frequency of 88MHz, the peak speed of a chip is 30.1 Gflops. Here, we follow the convention of assigning 38 operations for the calculation of pairwise gravitational force, which is adopted in resent Gordon-Bell prize applications. GRAPE-6 calculates the time derivative, which adds another 19 operations. Thus, the total number of floating point operations for one interaction is 57.

Each pipeline processor of GRAPE-6 calculates the force on eight particles simultaneously, using the virtual multiple pipeline (VMP) architecture introduced with GRAPE-4. In addition, six pipelines in one chip calculate the forces on six different sets of eight particles. Thus, one chip calculates the force on 48 particles. This architecture reduces the required memory bandwidth per chip drastically. The data stored for one particle is 64 bytes. Therefore, if six pipelines require different data, one chip requires 384 bytes per cycle, or 33.8GB/sec of the memory bandwidth. In our architecture, the bandwidth requirement is reduced by a factor of 48 to 0.7 GB/s, since the same data is reused 48 times to calculate force on 48 particles.

All processor chips in a cluster calculate the force on the same set of 48 particles, from different sets of particles. Thus, we need a fast network to first broadcast the 48 particles to all processor chips, and then summing up the result calculated on 256 chips in a cluster. The broadcast in hardware is pretty simple to implement in hardware. The summation network is implemented using FPGA chips as node points. Each FPGA chip implements a sequencer and 4-input adder. This FPGA and other circuits on board operate at 1/4 of

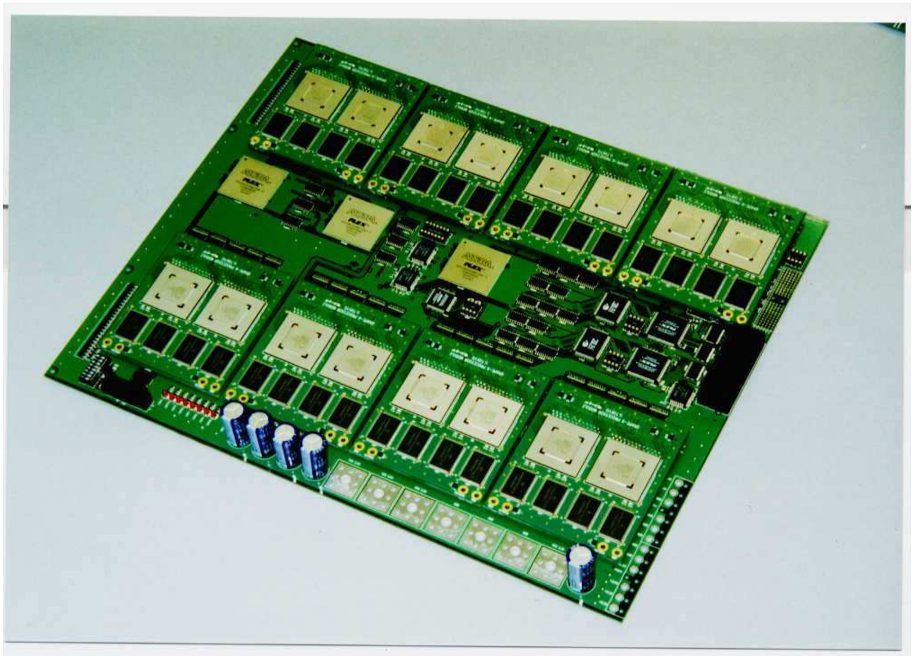

Figure 3: The GRAPE-6 processor board.

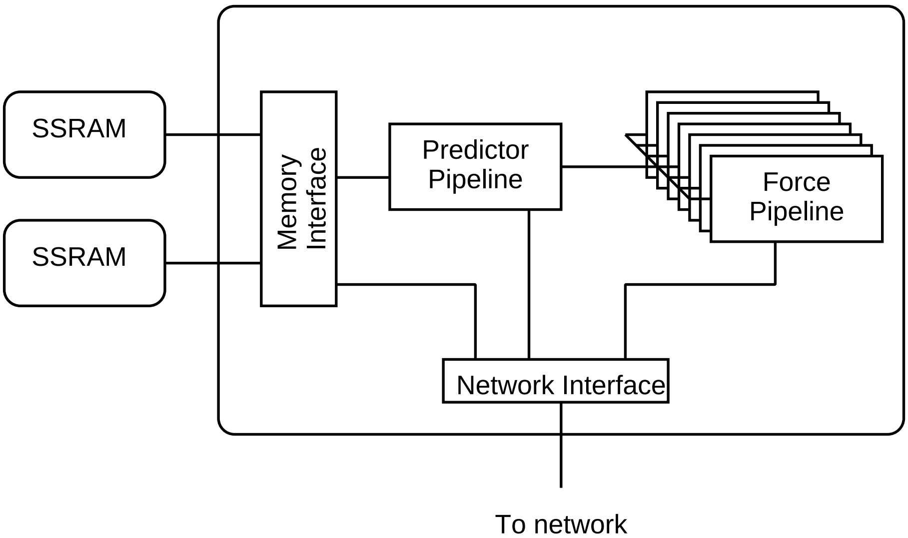

Figure 4: The GRAPE-6 processor chip.

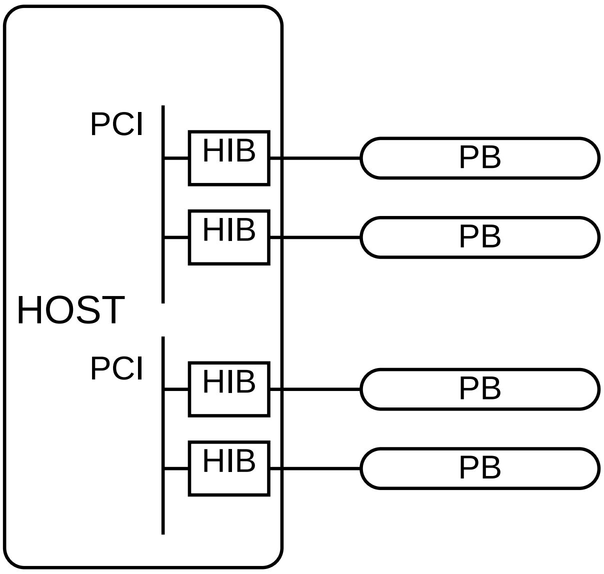

Figure 5: The present testbed GRAPE-6 configuration used for the simulation.

the clock speed of the chip. The data width for the network on board is 32 bits. So the communication speed of any link in the network is 88MB/s.

For the link between boards, we have adopted a fast semi-serial link with LVDS signal level, which can achieve the above data rate with four pairs of twisted-pair cables (the same as the standard cable for 100Mbit Ethernet). We adopted DS90CF364AMTD and DS90C363AMTD from National Semiconductor as the LVDS devices.

The structure of a NB is essentially the same as that of the processor board, but it carries links to the next level of the tree (either NB or PB) instead of the processor chips.

2.2 Development status

As of the time of writing, we have a 4-PB system, where each processor board is directly connected to the host computer without using the network board. Figure 5 shows the structure of this system. This system has a theoretical peak speed of 1.926 Tflops. The host computer is a UP-2000 based Alpha (EV6, 667 MHz) box with two processors and 2 GB of memory, running Compaq Tru64 UNIX.

3 Simulation and Performance

We have simulated the evolution of a galactic nucleus containing triple massive black holes. As we stated in the introduction, currently available numerical simulations suggest that when two galaxies merge, their central black holes will be left as a binary system. Thus, when the remnant of a galaxy merger again merges with a third galaxy, the final merged object is likely to host three massive black holes.

The evolution of a triple black hole system is entirely different from that of a binary system, since it is very unlikely that three black holes form a dynamically stable system. There are basically two possibilities for the final outcome. One is that one or more black holes will be ejected from the galaxy through three-body interaction of black holes. The other possibility is that during the three-body interaction two black holes come close enough that they merge through gravitational wave radiation. Thus, even after black holes have kicked out all nearby stars, they can still evolve in a complex way. On the other hand, a bound pair of two black holes is always dynamically stable. Thus, once a binary has kicked out all nearby stars, the evolution would be completely halted.

In our simulation, the galaxy is modeled with 786,432 equal-mass stars. The black holes are modeled as three point-mass particles with a mass of 1% of the total mass of the system. The relativistic effects were negligible in our simulation.

We performed a simulation for 2.5 dynamical time units, for which the number of individual steps was 1.549×10^9. The whole simulation, including on-the-fly analysis of the orbital elements of black holes and file operations, took 19.45 hours. The total number of floating point operations is $1.549 \times 10^9 \times 786,434 \times 57 = 6.944 \times 10^{16}$, since one particle-particle interaction amounts to 57 floating point operations. The resulting average computing speed is 992 Gflops.

As a short benchmark run, we tried the same simulation with 1,000,000 stars for 0.5 time units. For this run, the total number of timesteps is 2.407×10^8, and calculation took 3.43 hours. The average calculation speed was 1.11 Tflops.

4 Summary and Discussion

In this extended abstract, we present the performance achieved with astrophysical N-body simulations with individual timestep and direct force calculation on the GRAPE-6 special-purpose computer. The achieved performance number is 992 Gflops for a long calculation with 768k particles, and 1.11 Tflops for a short test calculation with 1 million particles.

The number of particles used is fairly large. Therefore, fast algorithms such as the Barnes-Hut tree algorithm might in principle outperform the direct summation algorithm we adopted here. However, as we've discussed in the introduction, there is no massively-parallel implementation of tree algorithm and individual timestep with accuracy high enough for this kind of problem. Therefore, we believe the speed we achieved is the fastest one can do with currently available combinations of hardware and software.

We thank Piet Hut for reading the manuscript. This work is supported by the Research for the Future Program of Japan Society for the Promotion of Science (JSPS-RFTF97P01102).

References

[Aar63] Aarseth Sverre J. (1963) Dynamical evolution of clusters of galaxies, i. *MNRAS*, 126: 223–255.

[EOG99] Eckart A., Ott T., and Genzel R. (December 1999) The Agr A* stellar cluster: New NIR imaging and spectroscopy. *A&A*, 352: L22–L25.

[JCR94] Jafry Y. R., Cornelisse J., and Reinhard R. (1994) LISA - A laser interferometer space antenna for gravitational-wave measurements. *ESA Journal*, 18: 219–228.

[MA93] McMillan S. L. W. and Aarseth S. J. (1993) An O(N log N) integration scheme for collisional stellar systems. *ApJ*, 414: 200–212.

[Mak97] Makino J. (1997) Merging of Galaxies with Central Black Holes. II. Evolution of the Black Hole Binary and the Structure of the Core. *ApJ*, 478: 58+.

[ME96] Makino J. and Ebisuzaki T. (1996) Merging of Galaxies with Central Black Holes. I. Hierarchical Mergings of Equal-Mass Galaxies. *ApJ*, 465: 527–533.

[MMH+95] Miyoshi M., Moran J., Herrnstein J., Greenhill L., Nakai N., Diamond P., and Inoue M. (January 1995) Evidence for a Black-Hole from High Rotation Velocities in a Sub-Parsec Region of NGC4258. *Nature*, 373: 127+.

[MT98] Makino J. and Taiji M. (1998) *Scientific Simulations with Special-Purpose Computers — The GRAPE Systems*. John Wiley and Sons, Chichester.

[MTES97] Makino J., Taiji M., Ebisuzaki T., and Sugimoto D. (1997) GRAPE-4: A Massively Parallel Special-Purpose Computer for Collisional N-Body Simulations. *ApJ*, 480: 432-446.

[QH97] Quinlan G. D. and Hernquist L. (November 1997) The dynamical evolution of massive black hole binaries - II. Self-consistent N-body integrations. *New Astronomy*, 2: 533–554.

[Qui96] Quinlan G. D. (1996) The dynamical evolution of massive black hole binaries I. Hardening in a fixed stellar background. *New Astronomy*, 1: 35–56.

[SCM+90] Sugimoto D., Chikada Y., Makino J., Ito T., Ebisuzaki T., and Umemura M. (1990) A special-purpose computer for gravitational many-body problems. *Nature*, 345: 33–35.

[TT72] Toomre A. and Toomre J. (December 1972) Galactic Bridges and Tails. *ApJ*, 178: 623–666.

[vv98] van der Marel R. P. and van den Bosch F. C. (November 1998) Evidence for a $3 \times 10^8 m_\odot$ Black Hole in NGC 7052 from Hubble Space Telescope Observations of the Nuclear Gas Disk. *AJ*, 116: 2220–2236.