

Special-purpose computing for dense stellar systems

Junichiro Makino

Division of Theoretical Astrophysics, National Astronomical Observatory of Japan,
2-21-1 Ohsawa, Mitaka, Tokyo 181-8588, Japan
email: makino@astron.s.u-tokyo.ac.jp

Abstract. I'll describe the current status of the GRAPE-DR project. The GRAPE-DR is the next-generation hardware for N -body simulation. Unlike the previous GRAPE hardwares, it is programmable SIMD machine with a large number of simple processors integrated into a single chip. The GRAPE-DR chip consists of 512 simple processors and operates at the clock speed of 500 MHz, delivering the theoretical peak speed of 512/226 Gflops (single/double precision). As of August 2006, the first prototype board with the sample chip successfully passed the test we prepared. The full GRAPE-DR system will consist of 4096 chips, reaching the theoretical peak speed of 2 Pflops.

Keywords. celestial mechanics, methods: n-body simulations, globular clusters.

1. Introduction

Since 1989, we have developed several GRAPE (GRAVity PipE) hardwares for astrophysical N -body simulations (Sugimoto *et al.* 1990; Makino & Taiji 1998). The first hardware, GRAPE-1 completed in 1989, had the speed of 240 Mflops. The present generation GRAPE-6 has the peak speed of 64 Tflops, and was the fastest computer at the time of its completion, 2002. We believe GRAPE project so far has been fairly successful, and we hope to continue to develop new hardwares.

A practical problem with a follow-on project for GRAPE-6 is that the initial cost of a custom chip has become too high. Initial cost of GRAPE-4 chip was USD 250 K. That of GRAPE-6 was around 1.5 M. A new chip will cost 5-20 M, depending on which company you talk to. The development cost of GRAPE-4 was already pretty large for a research project within the community of theoretical astrophysics of Japan. For GRAPE-6, we were very lucky to be selected as one project within a national program for 'Computational Science'. However, even with such a national program, USD 10 M just for the initial development of the chip which can be used only for astrophysical N -body simulation is way too much.

Therefore, we had to change our basic strategy of making highly specialized custom LSI chips for astrophysical N -body simulation. One option is to use FPGA (Field programmable Gate Array) chips to reduce the initial cost. This approach is quite effective if the required accuracy is not too high. However, even with the most advanced FPGA chips available today, the performance of high-accuracy force calculation is not much more than that of the 7-yr-old GRAPE-6 chip.

2. GRAPE-DR

Figure 1 shows the basic architecture of GRAPE-DR, the next-generation GRAPE hardware. It consists of a number of processing elements (PEs), each of which consists of

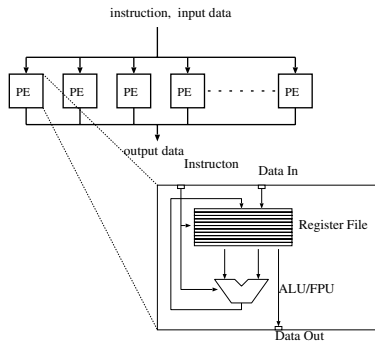


Figure 1. Basic structure of an SIMD processor

an FPU and a register file. They all receive the same instruction from outside the chip, and perform the same operation. This architecture is similar to that of classical SIMD machines such as Illiac-IV or TMC CM-2. Main difference is with GRAPE-DR we limit the size of the local memory of PEs so that we can fit a large number of PEs into a single chip.

With this architecture, each PE ‘emulates’ GRAPE pipeline by software. The absolute performance of such chip is around a factor of five less than full-custom GRAPE chip made with the same technology, because of additional transistors needed for memory and other control logics. However, with this approach we can vastly widen the application area, since PEs of GRAPE-DR are programmable by software. We hoped this wider application area would justify the large initial cost. We decided to call this architecture Greatly Reduced Array of Processor Elements, or GRAPE.

One practical problem with this new GRAPE architecture is that the number of PEs in one chip is too large. If each PE calculates the force on its own particle, a 512-PE GRAPE chip needs at least 512 particles to share the same time to attain good performance, when used with individual timestep algorithm. To reduce the number of PEs visible from the application program, we added a binary-tree reduction network which can add the results calculated on multiple PEs (Hence the name GRAPE-DR, with DR stands for Data Reduction). This tree has 16 inputs, and 512 processors are divided to 16 groups each with 32 PEs. These 32 PEs calculate force on different particles, and different groups calculate forces from different particles. This reduction tree turned to be useful for many other applications, including LINPACK. Sample chips arrived in May 2006 and operated successfully with the designed value of 500MHz clock on our prototype board. We are currently working on control logic on board, to run real applications.

Acknowledgements

This research is partially supported by the Special Coordination Fund for Promoting Science and Technology (GRAPE-DR project), Ministry of Education, Culture, Sports, Science and Technology, Japan.

References

- Makino, J., & Taiji, M. 1998, *Scientific Simulations with Special-Purpose Computers – The GRAPE Systems* (John Wiley and Sons)
- Sugimoto, D., Chikada, Y., Makino, J., Ito, T., Ebisuzaki, T., & Umemura, M. 1990, *Nature*, 345, 33