

A Workstation Farm Optimized for Monte Carlo Shell Model Calculations : Alphleet

Yasushi WATANABE,^{1,*} Noritaka SHIMIZU,^{1,2} Seigo HARUYAMA,² Michio HONMA,³
Takahiro MIZUSAKI,² Atsushi TAKETANI,¹
Yutaka UTSUNO² and Takaharu OTSUKA^{1,2}

¹*RIKEN, Saitama 351-0198, Japan*

²*Department of Physics, University of Tokyo, Tokyo 113-0033, Japan*

³*Center for Mathematical Sciences, University of Aizu, Fukushima 965-8580, Japan*

(Received December 9, 1999)

We have built a workstation farm named “Alphleet” which consists of 140 COMPAQ’s Alpha 21264 CPUs, for Monte Carlo Shell Model (MCSM) calculations. It has achieved more than 90 % scalable performance with 140 CPUs when the MCSM calculation with PVM and 61.2 Gflops of LINPACK.

The Quantum Monte Carlo Diagonalization (QMCD) method which is heart of the MCSM, has a peculiar characteristic as converting the original large matrix into a small one giving nearly equal eigenvalues by means of transforming basis vectors, as mentioned by another paper on this proceedings.¹⁾ In the QMCD, most important basis vectors of the many-body Hilbert space are searched in a stochastic way, and the Hamiltonian is diagonalized with respect to those basis vectors. Certainly, it takes almost all cpu time to calculate each Hamiltonian matrix element with many trial basis vectors to find out the (nearly) best transformation. This characteristic is suitable to implement on parallel calculation with using many CPUs, especially on loosely coupled system because of the calculation of each matrix element can be divided into individual segments not talking among themselves.

We decided to make such a computer system as known as “cluster” or “farm” by ourselves because in absence of the computer system which realize such simple parallelism in the market. Every parallel computer in the market is so sophisticate, that is, too expensive for our purpose. Performance of floating point calculation is also so crucial. Vector CPUs, state of the art architecture for fastest floating point calculation, are however not suitable for our calculation which does not treat vector long enough for the CPUs. We chose the Alpha 21264 which has highest performance of the floating point calculation (58.7 SPECfp95 at 500 MHz) as scaler CPU about 4 times better than most popular CPU in the commodity market.

The workstation farm that we have built is named “Alphleet”, which consists of many “Alpha” CPUs like as a “Fleet”. Figure 1 shows a schematic diagram and an actual photo (half nodes) of the Alphleet. It consists of a server and 70 clients.

It is our goal of the system to maximize performance for the MCSM calculation and to minimize management task. The latter means that easy management against such a large scale system is another big issue for non-professionals of computer itself

*) E-mail address: watanaby@riken.go.jp

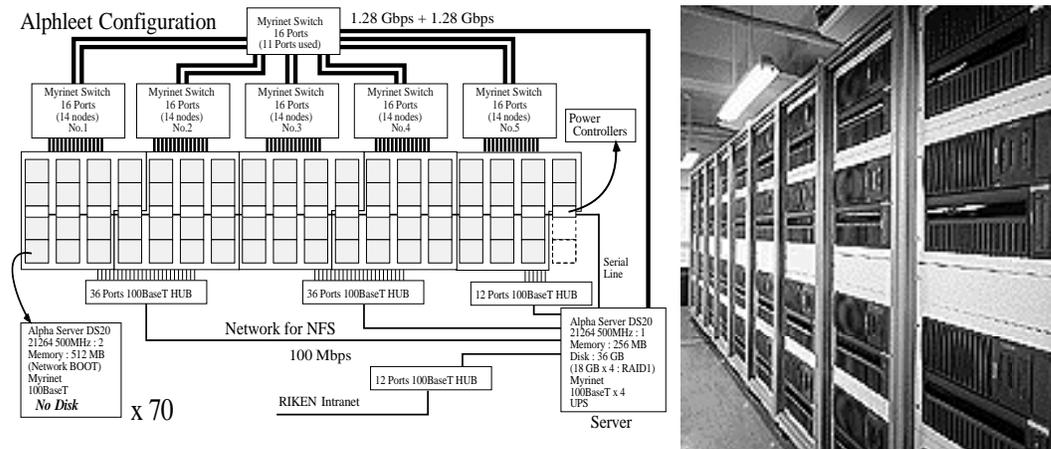


Fig. 1. A schematic diagram and a photo of the Alphleet.

like us. Therefore we decided to make disk-less configuration for clients with power controllers managed by the server. This means that each client has to be booted up by the server through network. The power controllers however make the server safe from overload when starting up the whole system with applying power on each client at intervals of 30 sec automatically. Thus the whole system is coming up about 40 min without any problem. Moreover, we can apply power on or off to each client with no care from remote site at any time, even a client was hung up without any response. No local swap space is another difficulty of disk-less configuration. So, we put 512 MB memory on each client which is estimated enough for the MCSM. This approach should not be only ours but more general solution for single purpose computer systems.

Each node has two network interfaces; one is 100Base-T for NFS used also for boot up and the other is Myrinet only for interprocess communications, which is 1.28 Gbps + 1.28 Gbps duplex switched network. Thanks to the Myrinet, the Alphleet achieved 93 % scalable performance when the MCSM calculation even on 140 CPUs. In addition, the Alphleet also achieved 61.2 Gflops of LINPACK benchmark without any precise tuning. It ranked in 169th on "14th TOP500 list" published at November 1999.

The Alphleet, which is built with many commodities, has carried out superior performance without managing difficulties. This is one example which has the advantage of single purpose system.

References

- 1) T. Otsuka, N. Shimizu, S. Haruyama, M. Honma, T. Mizusaki, A. Taketani, Y. Utsuno and Y. Watanabe, Prog. Theor. Phys. Suppl. No. 138 (2000), 24.