Performance Evaluation of Multiple Precision Numerical Computation using x86_64 Dual-core CPUs

Tomonori Kouya
tkouya@cs.sist.ac.jp
Shizuoka Institute of Science and Technology

1. Computing Environment

・Pentium D cluster: Pentium D 820 (2.8GHz), Fedora Core 4 x86-64, gcc-4.0.1, mpfr-2.2.0, gmp-4.1.4, MPICH2-1.0.2p, 2 node(max 4PEs)
・Pentium 4 cluster: Pentium IV 2.8cGHz, Vine Linux 3.1, gcc-3.4.3, mpfr-2.1.1, gmp-4.1.4, 11 node
・Xeon cluster: Xeon 3.0GHz (Dual Processor), Redhat 8, gcc-3.2, mpfr-2.1.1, gmp-4.1.4

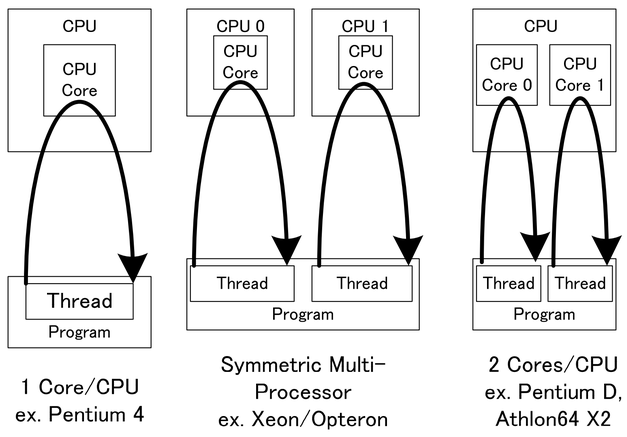

Fig.1: CPUs used in our experiments

2. GNU MP, MPFR, BNCpack and MPIBNCpack

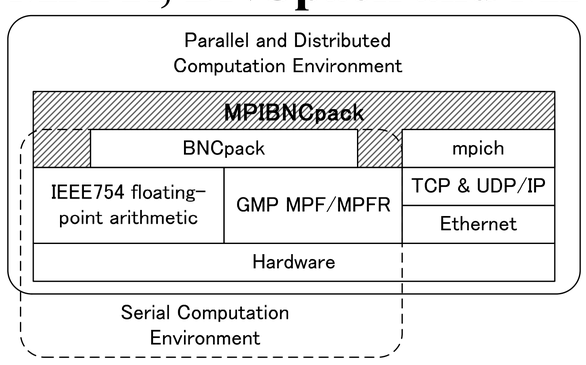

Fig.2 Software Structure of BNCpack/MPIBNCpack

BNCpack[3] is a Multiple Precision(MP) Numerical Computation Library based on GNU MP(GMP)[1] and MPFR[2]. MPIBNCpack [3,4] is a Parallelized MP Computation Library based on BNCpack and MPI. In 64bit environment, we can expect that GMP and MPFR will perform faster than 32bit environment because 64bit environment need the less number of operations per 1 MP arithmetic than 32bit one.

3. Performance Evaluation

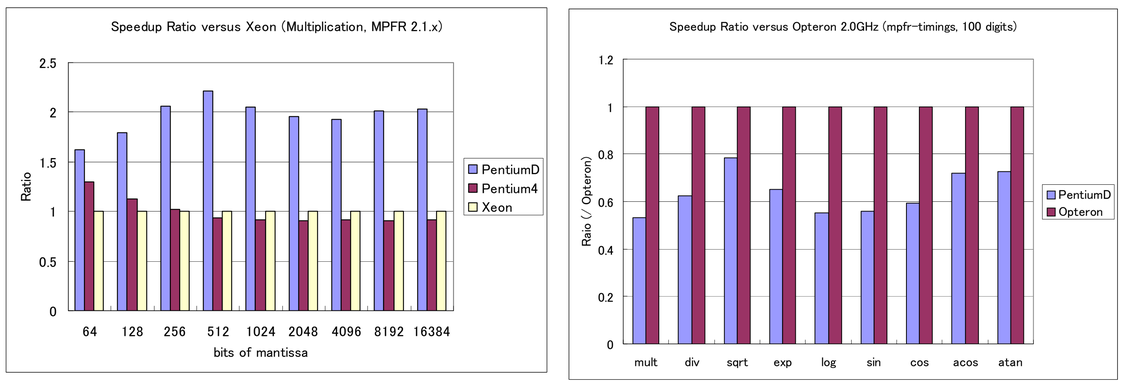

Fig. 3 Performance of Serial MP Computation
Left: vs Xeon, Right: vs Opteron

Our benchmark of serial MP arithmetic shows that the performance of Pentium D is about 1.6 to 2.2 times better than Pentium 4 and Xeon by using 64 bit operations. But the performance of Pentium D is about 0.5 to 0.8 times less than on 64bit environment with Opteron processor.

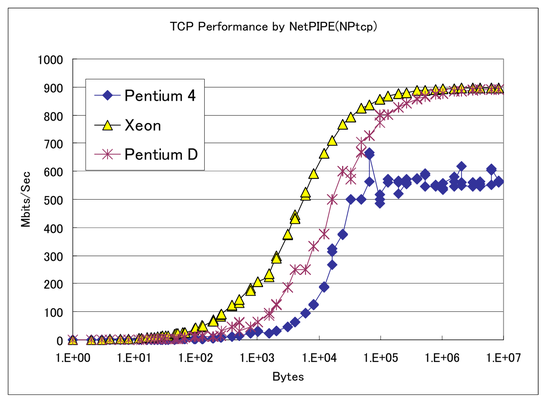

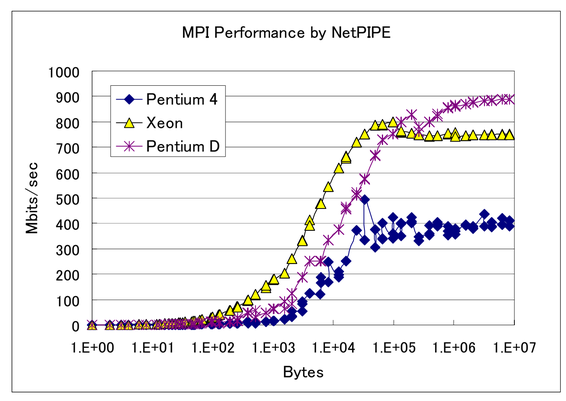

Fig. 4: Performance Comparison of TCP and MPI on GbE
Left: TCP , Right: MPI

Pentium 4, Xeon and Pentium D clusters which we can use are connected with 1000BASE-T Ethernet(GbE). Their performance comparison by NetPIPE[5] shows that our Pentium D cluster has the best performed GbE of all of them.

4. Performance of Parallel MP Computation

We have solved the following 512-dimensional linear system of equstion by CG method:

$$\begin{bmatrix} n & n-1 & \cdots & 1 \\ n-1 & n-1 & \cdots & 1 \\ \vdots & \vdots & & \vdots \\ 1 & 1 & \cdots & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{bmatrix}, x = \begin{bmatrix} 0 \\ 1 \\ \vdots \\ n-1 \end{bmatrix}$$

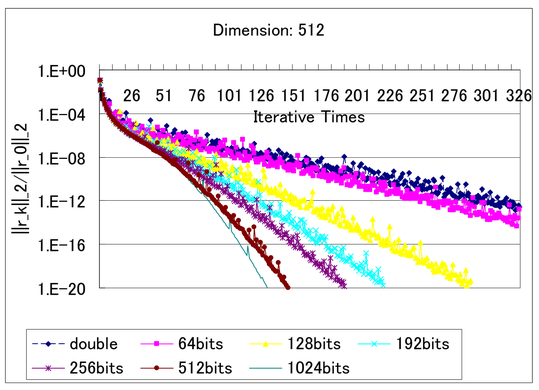

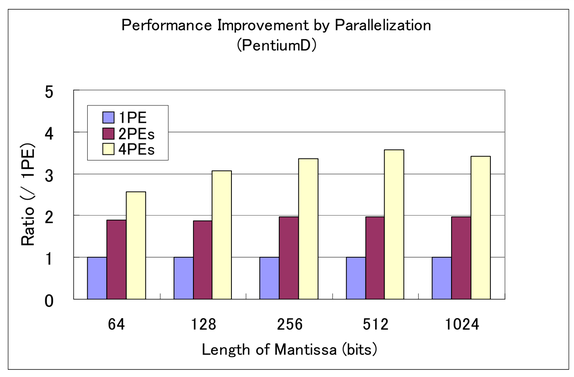

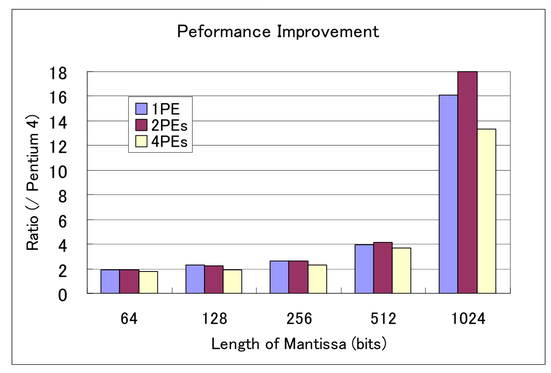

Fig. 5: CG Iteration Convergence History(Top) and Performance Improvement (Left: vs 1PE , Right: vs Pentium 4)

On our Pentium D cluster, we can confirm that:
・The parallel MP linear computation using 4PEs cannot obtain the best ideal performance we expected
・Their performance is better and better than Pentium 4 cluster as the length of mantissa is made longer.

References

[1] GNU MP, http://swox.com/gmp/
[2] MPFR Project, http://www.mpfr.org/
[3] BNCpack/MPIBNCpack, http://na-inet.jp/na/bnc/
[4] Tomonori Kouya, On BNCpack, a multiple precision numerical computation library, and its parallelization (in Japanese), Information, vol.7, no.4, pp.543-552, 2004.
[5] NetPIPE, http://www.scl.ameslab.gov/netpipe/