

FFTE: A High-Performance FFT Library and Parallel Implementation of Classical Gram-Schmidt Orthogonalization

Daisuke Takahashi

Center for Computational Sciences/
Graduate School of Systems and Information Engineering
University of Tsukuba

Outline

- FFTE: A High-Performance FFT Library
 - Features, Design and Approach
 - Performance Results
- Parallel Implementation of Classical Gram-Schmidt Orthogonalization Using Matrix Multiplication (with T. Yokozawa, T. Boku and M. Sato)
 - Column-wise Blocking CGS Algorithm
 - Recursive Blocking CGS Algorithm
 - Performance Results

FFTE: A High-Performance FFT Library

- FFTE is a Fortran subroutine library for computing the Fast Fourier Transform (FFT) in one or more dimensions.
- It includes complex, mixed-radix and parallel transforms.
- FFTE is typically faster than other publically-available FFT implementations, and is even competitive with vendor-tuned libraries.

FFTE Library - Features

- High speed
 - Supports Intel's SSE2/SSE3 instructions.
- Parallel transforms
 - Shared / Distributed memory parallel computers (OpenMP, MPI and OpenMP + MPI)
- High portability
 - Fortran77 + OpenMP + MPI
 - Intel's intrinsics for SSE2/SSE3 instructions.
- HPC Challenge Benchmark
 - FFTE's 1-D parallel FFT routine has been incorporated into the **HPC Challenge (HPCC) benchmark.**

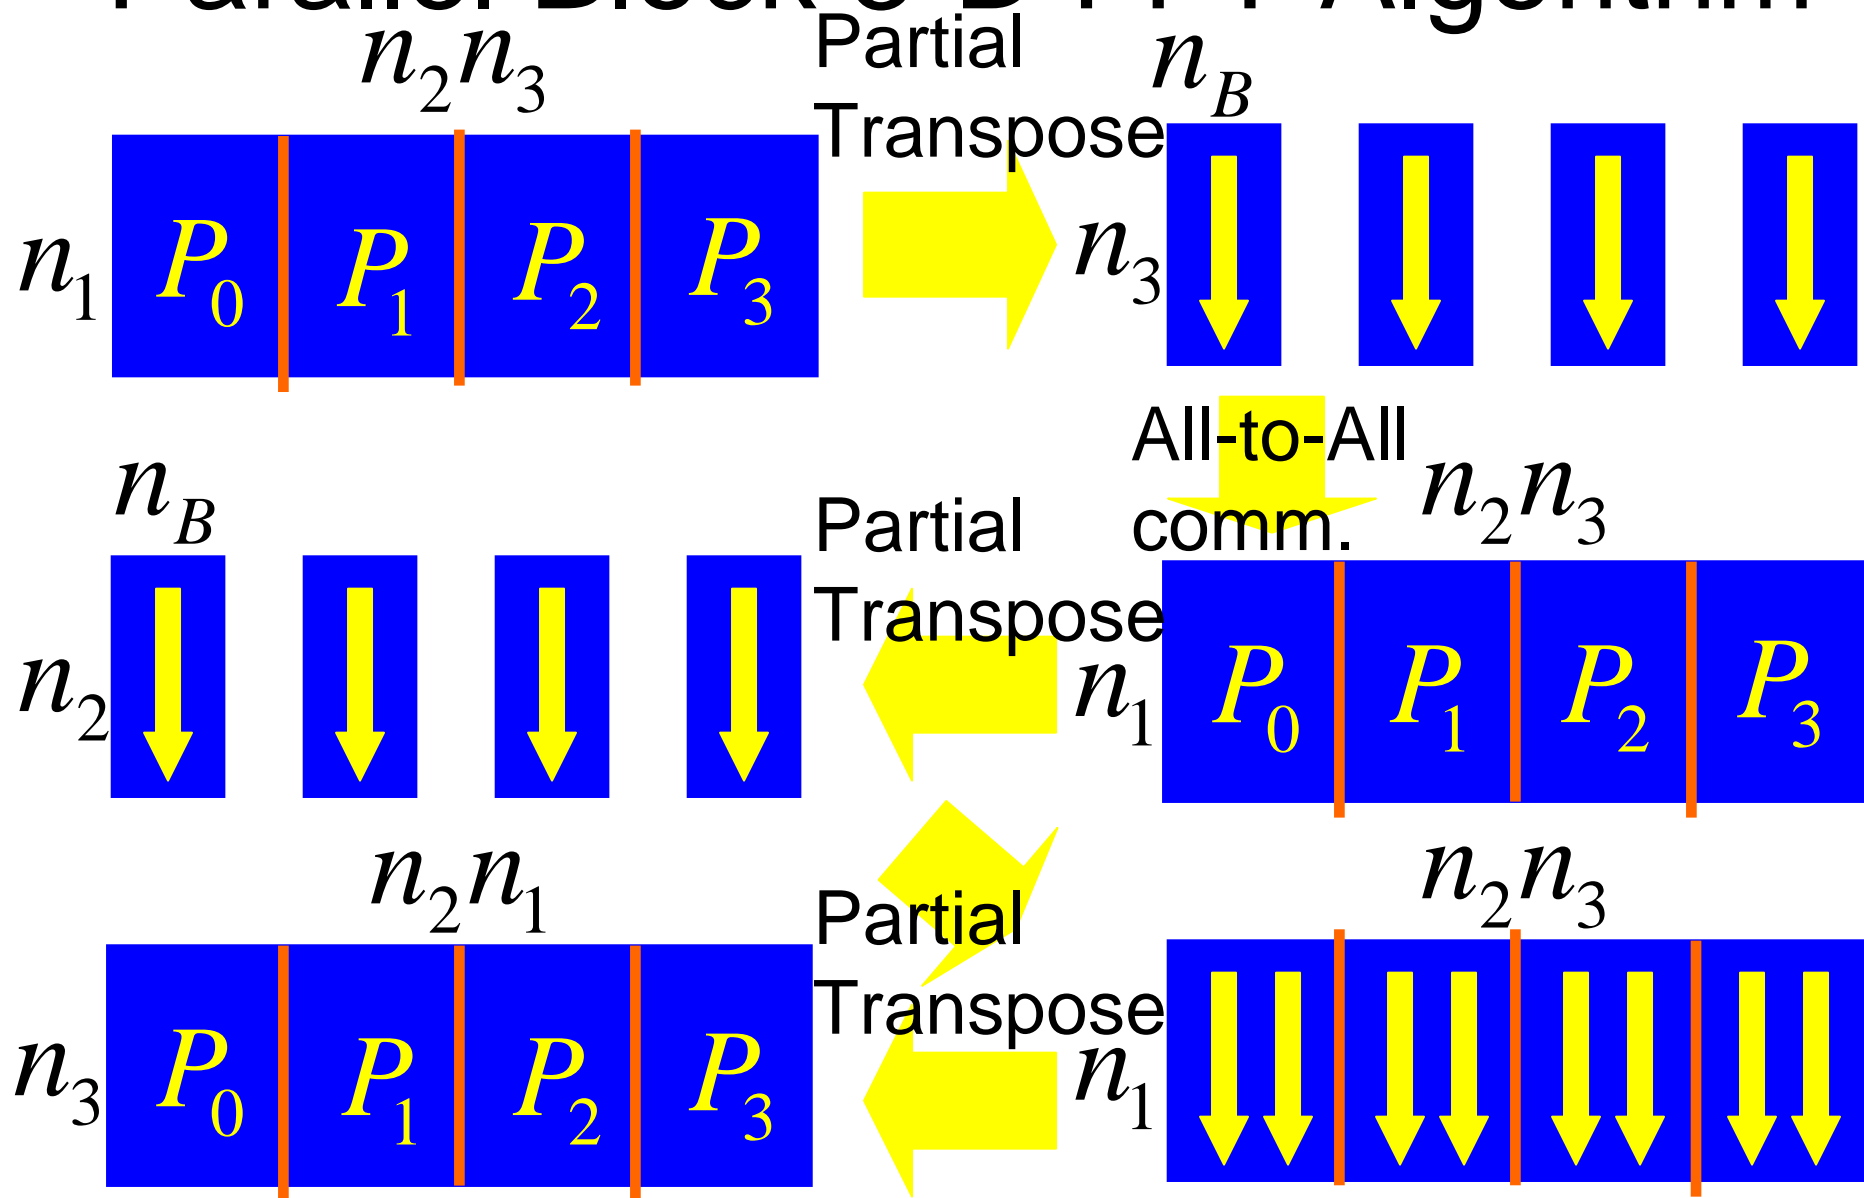
FFTE Library - Design

- Performance
 - One goal for large FFTs is to minimize the number of **cache misses**.
- Ease of use: routine interfaces
 - Similar to sequential **SGL SCSL** or **Intel MKL** routines
- Portability
 - Communication: MPI
 - Computation: Fortran77 + OpenMP

FFTE Library - Approach

- Many FFT routines work well when data sets **fit into a cache**.
- When a problem size exceeds the cache size, however, the performance of these FFT routines **decreases** dramatically.
- Some previously presented three-dimensional FFT algorithms require
 - Three multicolumn FFTs.
 - Three data **transpositions**.
 - The chief **bottlenecks** in cache-based processors.
- We combine the multicolumn FFTs and transpositions to **reduce** the number of cache misses.

Parallel Block 3-D FFT Algorithm




Performance Results

- To evaluate the implemented parallel 3-D FFTs, we compared
 - The implemented parallel 3-D FFT, named FFTE (ver 3.2, supports SSE2, using MPI)
 - FFTW (ver. 2.1.5, not support SSE2, using MPI)
- Target parallel machine:
 - A 16-node dual PC SMP cluster (dual Xeon 2.8GHz, 2GB DDR-SDRAM / node, Linux 2.4.20smp).
 - Interconnected through a Myrinet-2000 switch.
 - MPI-SCore [<http://www.pccluster.org>] was used.
 - We used an intra-node MPI library for the PC SMP cluster.


Performance of parallel 3-D FFTs

(dual Xeon PC SMP cluster, $N_1 \times N_2 \times N_3 = 2^{24} \times P$)



Breakdown of parallel 3-D FFTs

(dual Xeon PC SMP Cluster, $N_1 \times N_2 \times N_3 = 2^{24} \times P$)



Parallel Implementation of Classical Gram-Schmidt Orthogonalization

- Gram-Schmidt orthogonalization process is one of the fundamental algorithms in linear algebra.
 - Density-functional theory (DFT) code includes Gram-Schmidt orthogonalization of a large set of wave functions.
- Two basic computational variants of the Gram-Schmidt process exist:
 - Classical Gram-Schmidt algorithm (CGS)
 - Modified Gram-Schmidt algorithm (MGS)
 - Often selected for practical application
 - More stable than the CGS algorithm

Background

- Modified Gram-Schmidt (MGS) algorithm
 - Cannot be expressed by Level-2 BLAS.
 - Parallel implementation requires additional communication.
- Classical Gram-Schmidt (CGS) algorithm
 - Can be expressed by Level-2 BLAS.
 - CGS with DGKS correction [Daniel et al. '76] is one of the most efficient way to perform the orthogonalization process.

Classical Gram-Schmidt Algorithm

- The CGS orthogonalization can be performed by using Level-2 BLAS.
- Suitable for parallelization.
- The CGS algorithm is not stable.

do $j = 1, n$

$$\mathbf{q}_j = \mathbf{a}_j$$

do $i = 1, j - 1$

$$\mathbf{q}_j = \mathbf{q}_j - (\mathbf{q}_i, \mathbf{a}_j)\mathbf{q}_i$$

end do

$$\mathbf{q}_j = \mathbf{q}_j / \|\mathbf{q}_j\|$$

end do

\mathbf{a}_j Raw data vector

\mathbf{q}_j Orthogonalized vector

$\|\mathbf{q}_j\|$ Euclid norm

$(\mathbf{q}_i, \mathbf{a}_j)$ Inner product

Naive Implementation

- CGS for m by n matrix can be computed by
 - $(m - 1)$ matrix-vector multiplications (GEMV: Level-2 BLAS)
 - m normalizations (NRM2 and SCAL: Level-1 BLAS)

Let the matrix \mathbf{A} is denoted as $\mathbf{A} = (\mathbf{a}_1 \mathbf{a}_2 \cdots \mathbf{a}_n)$.

$$\mathbf{q}_1 = \mathbf{a}_1, \quad \underline{\mathbf{q}_1 = \mathbf{q}_1 / \|\mathbf{q}_1\|}$$

— GEMV

$$\underline{\mathbf{q}_2 = \mathbf{a}_2 - (\mathbf{q}_1, \mathbf{a}_2)\mathbf{q}_1}, \quad \underline{\mathbf{q}_2 = \mathbf{q}_2 / \|\mathbf{q}_2\|}$$

— NRM2, SCAL

$$\underline{\mathbf{q}_3 = \mathbf{a}_3 - (\mathbf{q}_1, \mathbf{a}_3)\mathbf{q}_1 - (\mathbf{q}_2, \mathbf{a}_3)\mathbf{q}_2}, \quad \underline{\mathbf{q}_3 = \mathbf{q}_3 / \|\mathbf{q}_3\|}$$

$$\underline{\mathbf{q}_4 = \mathbf{a}_4 - (\mathbf{q}_1, \mathbf{a}_4)\mathbf{q}_1 - (\mathbf{q}_2, \mathbf{a}_4)\mathbf{q}_2 - (\mathbf{q}_3, \mathbf{a}_4)\mathbf{q}_3}, \quad \underline{\mathbf{q}_4 = \mathbf{q}_4 / \|\mathbf{q}_4\|}$$

$$\underline{\mathbf{q}_5 = \mathbf{a}_5 - (\mathbf{q}_1, \mathbf{a}_5)\mathbf{q}_1 - (\mathbf{q}_2, \mathbf{a}_5)\mathbf{q}_2 - (\mathbf{q}_3, \mathbf{a}_5)\mathbf{q}_3 - (\mathbf{q}_4, \mathbf{a}_5)\mathbf{q}_4}, \quad \underline{\mathbf{q}_5 = \mathbf{q}_5 / \|\mathbf{q}_5\|}$$

$$\underline{\mathbf{q}_6 = \mathbf{a}_6 - (\mathbf{q}_1, \mathbf{a}_6)\mathbf{q}_1 - (\mathbf{q}_2, \mathbf{a}_6)\mathbf{q}_2 - (\mathbf{q}_3, \mathbf{a}_6)\mathbf{q}_3 - (\mathbf{q}_4, \mathbf{a}_6)\mathbf{q}_4 - (\mathbf{q}_5, \mathbf{a}_6)\mathbf{q}_5}, \quad \underline{\mathbf{q}_6 = \mathbf{q}_6 / \|\mathbf{q}_6\|}$$

CGS using Matrix Multiplication



- The CGS orthogonalization of a matrix can be changed into a matrix multiplication (GEMM: Level-3 BLAS).

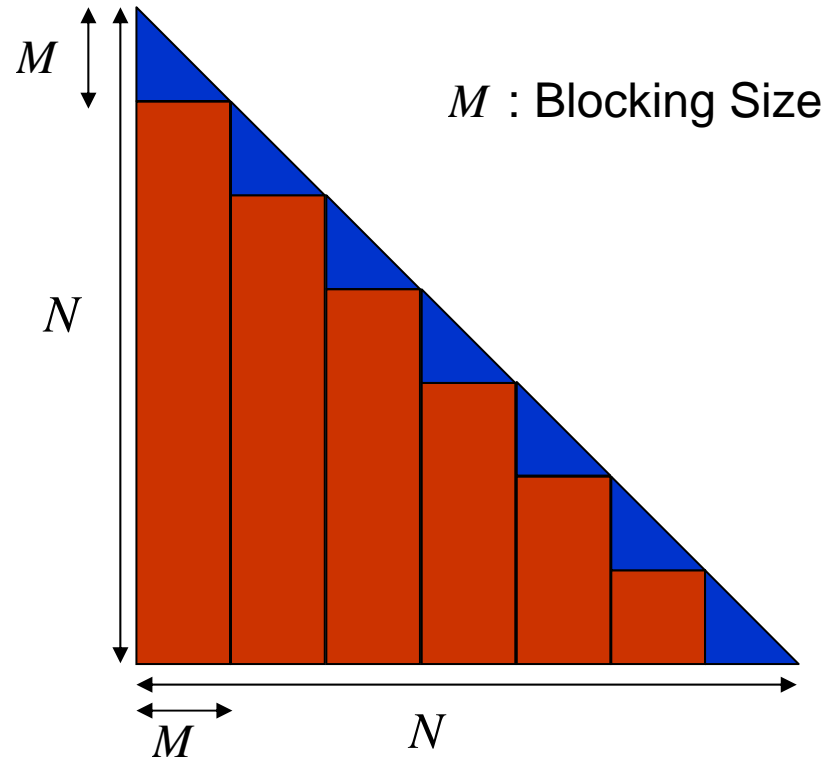
$$\begin{aligned} \mathbf{q}_1 &= \mathbf{a}_1, & \underline{\mathbf{q}_1} &= \underline{\mathbf{q}_1} / \|\mathbf{q}_1\| & & \text{--- GEMM} \\ \mathbf{q}_2 &= \mathbf{a}_2 - (\mathbf{q}_1, \mathbf{a}_2)\mathbf{q}_1, & \underline{\mathbf{q}_2} &= \underline{\mathbf{q}_2} / \|\mathbf{q}_2\| & & \text{--- GEMV} \\ \mathbf{q}_3 &= \mathbf{a}_3 - (\mathbf{q}_1, \mathbf{a}_3)\mathbf{q}_1 - (\mathbf{q}_2, \mathbf{a}_3)\mathbf{q}_2, & \underline{\mathbf{q}_3} &= \underline{\mathbf{q}_3} / \|\mathbf{q}_3\| & & \text{--- NRM2, SCAL} \\ \mathbf{q}_4 &= \mathbf{a}_4 - (\mathbf{q}_1, \mathbf{a}_4)\mathbf{q}_1 - (\mathbf{q}_2, \mathbf{a}_4)\mathbf{q}_2 - (\mathbf{q}_3, \mathbf{a}_4)\mathbf{q}_3, & \underline{\mathbf{q}_4} &= \underline{\mathbf{q}_4} / \|\mathbf{q}_4\| \\ \mathbf{q}_5 &= \mathbf{a}_5 - (\mathbf{q}_1, \mathbf{a}_5)\mathbf{q}_1 - (\mathbf{q}_2, \mathbf{a}_5)\mathbf{q}_2 - (\mathbf{q}_3, \mathbf{a}_5)\mathbf{q}_3 - (\mathbf{q}_4, \mathbf{a}_5)\mathbf{q}_4, & \underline{\mathbf{q}_5} &= \underline{\mathbf{q}_5} / \|\mathbf{q}_5\| \\ \mathbf{q}_6 &= \mathbf{a}_6 - (\mathbf{q}_1, \mathbf{a}_6)\mathbf{q}_1 - (\mathbf{q}_2, \mathbf{a}_6)\mathbf{q}_2 - (\mathbf{q}_3, \mathbf{a}_6)\mathbf{q}_3 - (\mathbf{q}_4, \mathbf{a}_6)\mathbf{q}_4 - (\mathbf{q}_5, \mathbf{a}_6)\mathbf{q}_5, & \underline{\mathbf{q}_6} &= \underline{\mathbf{q}_6} / \|\mathbf{q}_6\| \end{aligned}$$

Column-wise Blocking CGS

```

begin CBCGS(A, Q, N, M, s, h)
  do j = s, h, M
     $\mathbf{q}_j = \mathbf{q}_j / \|\mathbf{q}_j\|$ 
    do i = j + 1, j + M
       $\mathbf{w} = Q_{s,s+i}^T \mathbf{a}_{s+i}$ 
       $\mathbf{q}_{s+i} = Q_{s-i,s} \mathbf{w}$ 
       $\mathbf{q}_{s+i} = \mathbf{q}_{s+i} / \|\mathbf{q}_{s+i}\|$ 
    end do
     $W = Q_{j+M,N-(j+M)}^T A_{j,j+M}$ 
     $Q_{j+M,N-(j+M)} = WQ_{j+M,N-(j+M)}$ 
  end do
end
  
```

 Using Matrix-Vector Multiplication (GEMV)
 Using Matrix-Matrix Multiplication (GEMM)




Recursive Blocking CGS

```

begin RBCGS( $A, Q, N, M, s, h$ )
  if ( $h \leq M$ ) then
     $\mathbf{q}_s = \mathbf{q}_s / \|\mathbf{q}_s\|$ 
    do  $i = 1, h$ 
       $\mathbf{w} = Q_{s,s+i}^T \mathbf{a}_{s+i}$ 
       $\mathbf{q}_{s+i} = Q_{s-i,s} \mathbf{w}$ 
       $\mathbf{q}_{s+i} = \mathbf{q}_{s+i} / \|\mathbf{q}_{s+i}\|$ 
    end do
  else
    RBCGS( $A, Q, N, M, s, h/2$ )
     $W = Q_{s-(h/2),s+(h/2)}^T A_{s-(h/2),s+(h/2)}$ 
     $Q_{s-(h/2),s+(h/2)} = W Q_{s-(h/2),s+(h/2)}$ 
    RBCGS( $A, Q, N, M, s + h/2, h/2$ )
  end if
end

```

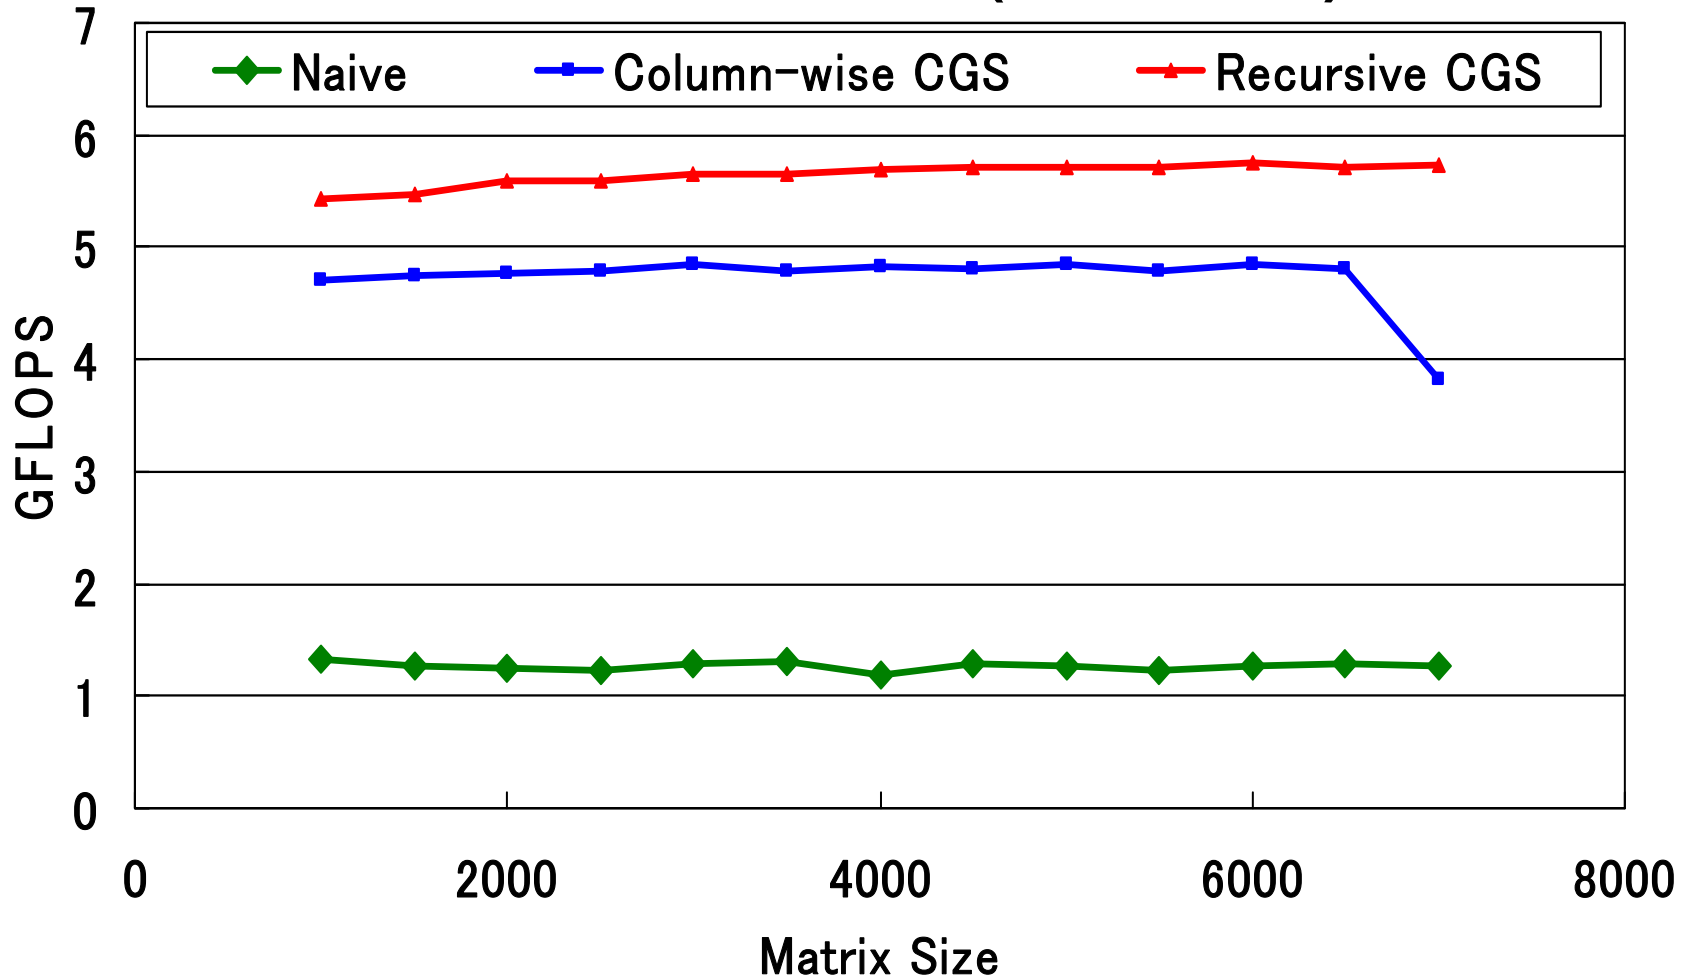
 Using Matrix-Vector Multiplication (GEMV)
 Using Matrix-Matrix Multiplication (GEMM)



Evaluation Environment

- A 32-node Xeon PC-cluster
 - Xeon 3GHz, 1GB DDR2 Memory
- 1000Base-T Gigabit Ethernet
- OpenMPI 1.2.3
- GotoBLAS r1.19
- Intel C Compiler 10.0
- Linux 2.6.20-1.2933.fc6smp
- All programs were run in 64-bit mode

Performance on Xeon 3GHz (1CPU)



Performance on 32 node 3GHz Xeon PC Cluster

