

Using Mixed Precision for Sparse Matrix Computations to Enhance the Performance while Achieving 64-bit Accuracy*

Alfredo Buttari¹, Jack Dongarra^{1,2}, Jakub Kurzak¹, Piotr Luszczek¹, and Stanimire Tomov¹

¹University of Tennessee Knoxville
²Oak Ridge National Laboratory

October 22, 2006

Abstract

By using a combination of 32-bit and 64-bit floating point arithmetic the performance of many sparse linear algebra algorithms can be significantly enhanced while maintaining the 64-bit accuracy of the resulting solution. These ideas can be applied to sparse multifrontal and supernodal direct techniques and sparse iterative techniques such as Krylov subspace methods. The approach presented here can apply not only to conventional processors but also to exotic technologies such as Field Programmable Gate Arrays (FPGA), Graphical Processing Units (GPU), and the Cell BE processor.

1 Introduction

Benchmarking analysis and architectural descriptions reveal that on many commodity processors the performance of 32-bit floating point arithmetic (single precision computations) may be significantly higher than 64-bit floating point arithmetic (double precision computations). This is due to a number of factors. First, many processors have vector instructions such as the SSE2 instruction set on the Intel IA-32 and IA-64 and AMD Opteron family of processors or the AltiVec unit on the IBM PowerPC architecture. In the SSE2 case a vector unit can complete four single precision operations every clock cycle but can complete only two in double precision. (For the AltiVec, single precision can complete eight floating point operations per cycle as opposed to

*This work was supported in part by the National Science Foundation and Department of Energy.

four floating point operations in double precision.) Another reason lies in the fact that single precision data can be moved at a faster rate through the memory hierarchy as a result of a reduced amount of data to be transferred. Finally, the fact that single precision data occupies less memory than double precision data also means that more single precision values can be held in cache than the double precision counterpart, which results in a lower rate of cache misses (the same reasoning can be applied to Translation Look-aside Buffers – TLBs). A combination of these factors can lead to significant enhancements in performance, as we will show for sparse matrix computations in the following sections. A remarkable example is the IBM Cell BE processor, where the peak performance using single precision floating point arithmetic is more than an order of magnitude higher than that of double precision (204.8 GFlop/s vs 14.6 GFlop/s for the 3.2 GHz version of the chip). The performance of some linear algebra operations can be improved based on the consideration that the most computationally expensive tasks can be performed by exploiting single precision operations and only resorting to double precision at critical stages, while attempting to provide the full double precision accuracy. This technique is supported by the well known theory of iterative refinement [13, 28], which has been successfully applied to the solution of dense linear systems [30]. This work is an extension of the work by Langou et al. [30] to the case of sparse linear systems, covering both direct and iterative solvers.

2 Sparse Direct and Iterative Solvers

Most sparse direct methods for solving linear systems of equations are variants of either multifrontal [18] or supernodal [5] factorization approaches. There is a number of freely available packages that implement these methods. We have chosen for our tests the software package MUMPS [1, 2, 3] as the representative of the multifrontal approach and SuperLU [32, 14, 15, 31] for the supernodal approach. Our main reason for selecting these two software packages is that they are implemented in both single and double precision, which is not the case for other freely available solvers such as UMFPACK [19, 11, 12].

Fill-ins, and the associated memory requirements, are inherent for direct sparse methods. And although there are various reordering techniques designed to minimize the amount of these fill-ins, for problems of increasing size there is a point where the memory requirements become prohibitively high and direct sparse methods are no longer feasible. Iterative methods are a remedy because only a few working vectors and the primary data are required [8, 37].

Two popular iterative solvers on which we will illustrate the techniques addressed in this paper are the Conjugate Gradient (CG) method (for symmetric and positive definite matrices) and the Generalized Minimal Residual (GMRES) method for non-symmetric matrices [38]. The preconditioned versions of the two algorithms are given correspondingly in Algorithms 1 and 2 below with the descriptions that follow the standard notation [8, 37].

The preconditioners, denoted in both cases as M , are operators intended to improve the robustness and the efficiency of the iterative algorithms. In particular, we will use

Algorithm 1 PCG (b, x_0, E_{tol}, \dots)

```
1:  $r_0 = b - Ax_0$ 
2: for  $i = 1$  to ... do
3:    $z_{i-1} = Mr_{i-1}$ 
4:    $\rho_{i-1} = r_{i-1}^T z_{i-1}$ 
5:   if  $i = 1$  then
6:      $d_1 = z_0$ 
7:   else
8:      $\beta = \rho_{i-1} / \rho_{i-2}$ 
9:      $d_i = z_{i-1} + \beta d_{i-1}$ 
10:  end if
11:   $q_i = Ad_i$ 
12:   $\alpha = \rho_{i-1} / d_i^T q_i$ 
13:   $x_i = x_{i-1} + \alpha d_i$ 
14:   $r_i = r_{i-1} - \alpha q_i$ 
15:  check convergence; continue if necessary
16: end for
```

Algorithm 2 GMRES ($b, x_0, E_{tol}, m, \dots$)

```
1: for  $i = 0, 1, \dots$  do
2:    $r = b - Ax_i$ 
3:    $\beta = h_{1,0} = \|r\|_2$ 
4:   check convergence and exit if done
5:   for  $k = 1, \dots, m$  do
6:      $v_k = r / h_{k,k-1}$ 
7:      $r = A M v_k$ 
8:     for  $j=1, \dots, k$  do
9:        $h_{j,k} = r^T v_j$ 
10:       $r = r - h_{j,k} v_j$ 
11:    end for
12:     $h_{k+1,k} = \|r\|_2$ 
13:    if  $h_{k+1,k}$  is small enough then break
14:  end for
15:  // Define  $V_k = [v_1, \dots, v_k]$ ,  $H_k = \{h_{i,j}\}_{1 \leq i \leq k+1, 1 \leq j \leq k}$ 
16:  Find  $W_k = [w_1, \dots, w_k]^T$  that minimizes  $\|b - A(x_i + M V_k W_k)\|_2$ 
17:  // note: or equivalently, find  $W_k$  that minimizes  $\|\beta e_1 - H_k W_k\|_2$ 
18:   $x_{i+1} = x_i + M V_k W_k$ 
19: end for
```

left preconditioning, where instead of

$$Ax = b$$

we solve $MAx = Mb$, and right preconditioning, where the problem is transformed to $AMu = b$, $x = Mu$. Intuitively, to serve its purpose, M needs to be easy to compute,

apply and store as well as to approximate A^{-1} .

The basic idea of our approach is to use faster but lower precision computations whenever possible. As we show in the rest of the paper, this idea can be used to design the preconditioner M that has two requirements mentioned above. And since our basic idea can be exploited (in iterative solvers) through proper preconditioning, the applicability of the approach is far-reaching and not limited to neither the preconditioners nor the solvers each of which are used to demonstrate our idea.

3 Mixed precision Iterative Refinement

The iterative refinement technique is a well known method that has been extensively studied and applied in the past. A fully detailed description of this method can be found elsewhere [13, 28, 43, 51, 9]. The iterative refinement approach has been used in the past to improve the accuracy of linear systems' solutions and it is shown in Algorithm 3.

Algorithm 3 The iterative refinement method for the solution of linear systems

```
1:  $x_0 \leftarrow A^{-1}b$ 
2:  $k = 1$ 
3: repeat
4:    $r_k \leftarrow b - Ax_{k-1}$ 
5:    $z_k \leftarrow A^{-1}r_k$ 
6:    $x_k \leftarrow x_{k-1} + z_k$ 
7:    $k \leftarrow k + 1$ 
8: until convergence
```

Once the system is solved at step 1, the solution can be refined through an iterative procedure where, at each iteration, the residual is computed based on the solution at the previous iteration (step 4), a correction is computed as in step 5, and finally this correction is applied as in step 6. While the common usage of iterative refinement [29, 4] consists of performing all the arithmetic operations with the same precision (either single or double), we have investigated the application of mixed precision, iterative refinement where the most expensive steps, 1 and 5, are performed in single precision and steps 4 and 6 are performed in double precision. We are not aware of related work [45, 44, 22] that uses our exact approach. The error analysis for the mixed precision, iterative refinement, explained in [33, 21, 23], shows that using this approach it is possible to achieve the same accuracy as if the system was solved in full double precision arithmetic, provided that the matrix is not too badly conditioned. From a performance point of view, the potential of this method lies in the fact that the most computationally expensive steps, 1 and 5, can be performed very fast in single precision arithmetic while the only tasks that require double precision accuracy are steps 4 and 6, whose cost can be considered much lower.

We will refer below to single precision (SP) as 32-bit and to double precision (DP) as 64-bit floating point arithmetic, and also lower and higher precision arithmetic will be correspondingly associated with SP and DP.

3.1 Mixed precision Iterative Refinement for Sparse Direct Solvers

Using the MUMPS package for solving systems of linear equations can be described in three distinct steps:

1. System Analysis: in this phase the system sparsity structure is analyzed in order to estimate the fill-in, which provides an estimate of the memory requirement that will be allocated in the following steps. Also, pivoting is performed based on the structure of $A + A^T$, ignoring numerical values. Only integer operations are performed at this step.
2. Matrix Factorization: in this phase, the $PA = LU$ factorization is performed. This is the computationally most expensive step of the system solution.
3. System Solution: the system is solved in two steps: $Ly = Pb$ and $Ux = y$.

Once steps 1 and 2 are performed, each iteration of the refinement loop needs only to perform the system solution (i.e., step 3), whose cost can be up to two orders of magnitude lower than the cost of the system factorization. The implementation of the mixed precision, iterative refinement method with the MUMPS package is shown in Algorithm 4.

Algorithm 4 mixed precision Iterative Refinement with the MUMPS package

- 1: system analysis
 - 2: $LU \leftarrow PA$ (SP)
 - 3: solve $Ly = Pb$ (SP)
 - 4: solve $Ux_0 = y$ (SP)
 - 5: **until** convergence **do**:
 - 6: $r_k \leftarrow b - Ax_{k-1}$ (DP)
 - 7: solve $Ly = Pr_k$ (SP)
 - 8: solve $Uz_k = y$ (SP)
 - 9: $x_k \leftarrow x_{k-1} + z_k$ (DP)
 - 10: **done**
-

At the end of each line of the algorithm, we indicate the precision used to perform the corresponding operation. Based on backward stability analysis, we consider that the solution x is of double precision quality when

$$\|b - Ax\|_2 \leq \|x\|_2 \cdot \|A\|_{fro} \cdot \epsilon \cdot \sqrt{n}$$

where $\|\cdot\|_{fro}$ is the Frobenius norm and n is the problem size. This provides us a stopping criterion. If some maximum number of iterations is reached, then the algorithm should signal failure to converge. All the control parameters for the MUMPS solver have been set to their default values, which means that the matrix scaling and permuting and pivoting order strategies are determined at runtime based on the matrix properties.

3.2 Mixed precision Iterative Refinement for Sparse Iterative Solvers

The general framework of mixed precision, iterative refinement given at the beginning of this section can be easily extended to sparse iterative solvers. Indeed, the mixed precision, iterative refinement can be interpreted as a preconditioned Richardson iteration with the preconditioner computed and applied (during the iterations) in single precision [46]. This interpretation can be further extended from Richardson to any preconditioned, iterative method. And in general, as long as the iterative method at hand is backward stable and converges, one can apply similar reasoning as in [30] to show that the solution obtained would be accurate in higher precision. The feasibility of using mixed precision in computing and/or applying a preconditioner depends first on whether there is a potential to introduce speedups in the computation, and second on how the method's robustness would change.

A way to use mixed precision in the preconditioner is when, for example, any higher precision data is simply replaced with a corresponding lower precision, but all the computations are still done in higher precision. The success of this approach, regarding speed, depends on what percent of the overall computation is spent on the preconditioner. For example, a simple diagonal preconditioner may not benefit from it, while a domain decomposition-based [35], block diagonal preconditioner, or a multigrid V-cycle [27], may benefit. Also, multigrid-based solvers may benefit both in speed (as the bulk of the computation is in their V/W-cycles) and memory requirements. An example of successful application of this type of approach in CFD was done [25, 26] in a PETSc solver, which was accelerated with a Schwartz preconditioner using block-incomplete factorizations over the separate subdomains that are stored in single precision. Regarding robustness, there are various algorithmic issues to consider, including ways to automatically, at run time, determine limitations of the approach. This brings up another possible idea to explore, which is the use of lower precision arithmetic only for parts of the preconditioner. Examples here may come from adaptive methods that automatically locate the singularities of the sought solution, and hence the corresponding parts of the matrix responsible for resolving them. This information may be used in combination with the solver and preconditioner (e.g., hierarchical multigrid) to achieve both speedup and robustness of the method.

Another interesting example is when not just the higher precision storage, but also the higher precision arithmetic are replaced with lower precision. This is the case that would allow one to apply the technique not only to conventional processors, but also to FPGAs, GPUs, Cell BE processor, etc.

The focus of the current work is to enable the efficient use of lower precision arithmetic to sparse iterative methods in general, when no preconditioner or when just a simple and computationally inexpensive (relative to the rest of the computation) preconditioner is available. The idea of accomplishing this is to use the preconditioned version of the iterative method at hand and replace the preconditioner M by an iterative method as well but implemented in reduced precision arithmetic. Thus, by controlling the accuracy of this iterative, *inner* solver, more computations can be done in reduced precision and less work is needed in the full precision arithmetic.

The robustness of variations of this *nesting* of iterative methods, also known in the

literature as *inner-outer* iteration, has been studied before, both theoretically and computationally [24, 36, 41, 6, 34, 50, 48]. The general appeal of these methods is that computational speedup is possible when the inner solver uses an approximation to the original matrix that is also faster to apply. Moreover, even if no faster matrix-vector product is available, speedup can often be observed due to improved convergence (e.g., see restarted GMRES vs GMRES-FGMRES [41] and Subsection 4.3). To our knowledge, using mixed precision for performance enhancement has not been done in the framework suggested in this paper. In the subsections below, we show a way to do it for CG and GMRES.

3.2.1 CG-based Inner-Outer Iteration Methods

We suggest the PCG-PCG inner-outer Algorithm 5. It is given as a modification to the reference PCG Algorithm 1, and therefore only the lines that change are written out (i.e., line 3). The inner PCG, denoted by PCG_single, is in SP and is described by Algorithm 6 again as a modification of the reference PCG. The preconditioner available for the reference PCG is used in SP in the inner PCG. Note that our initial guess in the inner loop is always taken to be 0 and we perform a fixed number of iterations (step 15 is in brackets since practically we want to avoid exiting due to small residual, unless it is of the order of the machine’s single precision). The number of inner iterations is fixed and depends on the particular problem at hand. We take it to be the number of iterations it takes PCG_single to do a fixed (e.g., 0.3) relative reduction for the initial residual r_0 . Work on criteria to compute the (variable) number of inner iterations guaranteeing convergence can be found in [40].

Algorithm 5 PCG_PCG (b, x_0, E_{tol}, \dots)

```

...
3: PCG_single (  $r_{i-1}, z_{i-1}, NumIters, \dots$  )
...

```

Algorithm 6 PCG_single ($b, x, NumIters, \dots$)

```

1:  $r_0 = b; x_0 = 0$ 
2: for  $i = 1$  to  $NumIters$  do
...
15: [check convergence and break if done]
16: end for

```

If a preconditioner is not available, we can similarly define a CG-PCG algorithm, where the inner loop is just a CG in SP. Furthermore, other iterative solvers can be used for the inner loop, as long as they result in symmetric and positive definite (SPD) operators. For example, stationary methods like Jacobi, Gauss-Seidel (combination of one backward and one forward to result in SPD operator), and SSOR can be used.

Note that with these methods, a constant number of iterations and initial guess 0 result in a constant preconditioner and hence in optimal convergence for the outer PCG iteration. The use of a Krylov space method in the inner iteration, as in the currently considered algorithm, results in a non-constant preconditioner. Although there is convergence theory for these cases [41], how to set the stopping criteria still remains to be resolved as do variations in the algorithms, etc. [24, 34] in order to obtain optimal results. For example, G. Golub and Q. Ye [24] consider the inexact PCG (β is taken as $\frac{(r_{i-1}-r_{i-2})z_{i-1}}{r_{i-2}z_{i-2}}$), which allows certain local orthogonality relations to be preserved from the standard PCG, which on the other hand provides grounds for theoretically studied aspects of the algorithm. We tried this approach as well and, although our numerical results were similar to [24], overall the algorithm described here gave better results. In general, non-constant preconditioning deteriorates the CG convergence, often resulting in convergence that is characteristic of the Steepest Descent algorithm. Still, shifting the computational load to the inner PCG reduces this effect and gives convergence that is comparable to the convergence of a reference PCG algorithm. Note that in our case, by setting an inner number of iterations so that a fixed relative reduction for the residual is achieved, we expect to have only a constant number of outer iterations until convergence. We note that non-constant preconditioning can be better accommodated in GMRES (next section). See also Simoncini and Szyld [42] for a way to interpret and theoretically study the effects of non-constant preconditioning.

3.2.2 GMRES-based Inner-Outer Iteration Methods

For our outer loop, we take the flexible GMRES (FGMRES [36, 37]). This is a minor modification to Algorithm 2 meant to accommodate non-constant preconditioners. Note that on line 7, the preconditioner M is non-constant and actually depends on v_k , so the update of x_i at line 18 will make the algorithm unstable if first $V_m W_m$ is computed and then M applied. This can be remedied for the price of m additional storage vectors. Algorithm 7 gives our inner-outer GMRES-FGMRES (the additional vectors are introduced at line 7: $z_k = Mv_k$ where M is replaced with the `GMRES_single` solver in Algorithm 8).

Algorithm 7 GMRES_FGMRES ($b, x_o, E_{tol}, m_{in}, m_{out}, \dots$)

```

...
5: for  $k = 1$  to  $m_{out}$  do
...
7:   GMRES_single(  $r_k, z_k, 1, m_{in}, \dots$  )
    $r = Az_k$ 
...
18:   $x_{i+1} = x_i + Z_k W_k$ 
19: end for

```

As with PCG-PCG, the algorithm is given as a modification to the reference GMRES(m) algorithm from Algorithm 2 and, therefore, only the lines that change are

Algorithm 8 GMRES_single ($b, x, NumIters, m, \dots$)

```
1:  $x_0 = 0$ 
   for  $i = 1$  to  $NumIters$  do
     ...
4:   [check convergence and break if done]
     ...
```

written out. The inner GMRES_single is in SP. The preconditioner available for the reference GMRES is used in SP in the inner GMRES_single. Note that again our initial guess in the inner loop is always taken to be 0 and we perform a fixed number of cycles/restarts (in this case just 1; step 4 is in brackets since we want to avoid exiting due to small residual, unless it is of the order of the machine's single precision).

The potential benefits of FGMRES compared to GMRES are becoming better understood [41]. Numerical experiments, as we also show, confirm cases of improvements in speed, robustness, and sometime memory requirements for these methods. For example, we show a maximum speedup of close to $12\times$ on a problem of size 602,091 (see Section 4). The memory requirements for the method are the matrix in CRS format, the nonzero matrix coefficients in SP, twice the outer restart size number of vectors in DP, and inner restart size number of vectors in SP.

The Generalized Conjugate Residuals (GCR) method [50, 49] is comparable to the FGMRES and can replace it successfully as the outer iterative solver.

4 Results Overview

4.1 The Test Collection for Mixed precision Sparse Direct and Iterative Solvers

We tested our implementation of a mixed precision, sparse direct solver on a test suite of 41 matrices taken from the University of Florida's Sparse Matrix Collection [47]. The matrices were selected randomly from the collection since there is no information available about their condition number. A smaller subset of nine matrices (described in Table 1) will be discussed in this document for readability reasons. The matrices in this smaller subset were chosen in order to provide examples of all the significant features observed on the test suite. The results for all the 41 matrices in the test suite can be found in [10].

For the iterative sparse methods we used matrices from an adaptive 3D PDEs discretization package. Presented are results on a set of five matrices of increasing size, coming from the adaptive discretization of a 3D linear elasticity problem on a tetrahedral mesh, using piecewise linear elements (see Table 2).

Matrix No.	Matrix name	Size	Nonzeroes	Cond. num. est.
1	G64	7000	82918	$O(10^4)$
2	Si10H16	17077	875923	$O(10^3)$
3	c-71	76638	859554	$O(10)$
4	cage11	39082	559722	$O(1)$
5	dawson5	51537	1010777	$O(10^4)$
6	nasasrb	54870	2677324	$O(10^7)$
7	poisson3Db	85623	2374949	$O(10^3)$
8	rma10	46835	2374001	$O(10)$
9	wang4	26068	177196	$O(10^3)$

Table 1: Properties of a subset of the tested matrices. The condition number estimates were computed on the Opteron 246 architecture by means of MUMPS subroutines.

Level	Size	Nonzeroes	Cond. num. est.
1	11,142	442,225	$O(10^3)$
2	25,980	1,061,542	$O(10^4)$
3	79,275	3,374,736	$O(10^4)$
4	230,793	9,991,028	$O(10^5)$
5	602,091	26,411,323	$O(10^5)$

Table 2: Properties of the matrices used with the iterative sparse solvers.

4.2 Performance Characteristics of the Tested Hardware Platforms

The implementation of the mixed precision algorithm for sparse direct methods presented in Section 3.1 has been tested on the architectures reported, along with their main characteristics, in Table 3. All the tests were done with sequential code; thus only one execution unit was used even on those processors that present multiple cores. All of these architectures have vector units except the Sun UltraSparc-IIe one; this architecture has been included with the purpose of showing that even in the case where the same number of single and double precision operations can be completed in one clock cycle, significant benefits can still be achieved thanks to the reduced memory traffic and higher cache hit rate provided by single precision arithmetic.

The implementation of the mixed precision algorithms for sparse iterative solvers described in Section 3.2 was only tested on the Intel Woodcrest architecture. The application for the numerical tests on the mixed precision iterative method for sparse direct solvers was coded in Fortran 90 and the application for the numerical tests on the inner-outer iteration method for sparse iterative solvers was coded in C.

Table 4 shows the difference in performance between the single and double precision implementation for the two dense BLAS operations matrix-matrix product (`_GEMM`) and matrix-vector product (`_GEMV`). They are the two principal computational kernels of sparse direct solvers: sparse data structures get rearranged to fit

	Clock freq.	Vector Units	Memory	Compiler	Compiler flags	BLAS
AMD Opteron 246	2 GHz	SSE, SSE2 3DNOW!	2 GB	Intel v9.1	-O3 -fast	Goto
Sun UltraSparc-IIe	502 MHz	none	512 MB	Sun v9.0	-xchip=ultra2e -xarch=v8plusa	Sunperf
Intel PIII Copp.	900 MHz	SSE, MMX	512 MB	Intel v9.0	-O3	Goto
PowerPC 970	2.5 GHz	AltiVec	2 GB	IBM v8.1	-O3 -qalign=4k	Goto
Intel Woodcrest	3 GHz	SSE, SSE2 MMX	4 GB	Intel v9.1	-O3	Goto
Intel XEON	2.4 GHz	SSE, SSE2 MMX	2 GB	Intel v8.0	-O3	Goto
Intel Centrino Duo	2.5 GHz	SSE, SSE2 MMX	4 GB	Intel v9.0	-O3	Goto

Table 3: Characteristics of the architectures used to measure the experimental results.

	Size	SGEMM/ DGEMM	Size	SGEMV/ DGEMV
AMD Opteron 246	3000	2.00	5000	1.70
Sun UltraSparc-IIe	3000	1.64	5000	1.66
Intel PIII Copp.	3000	2.03	5000	2.09
PowerPC 970	3000	2.04	5000	1.44
Intel Woodcrest	3000	1.81	5000	2.18
Intel XEON	3000	2.04	5000	1.82
Intel Centrino Duo	3000	2.71	5000	2.21

Table 4: Performance comparison between single and double precision arithmetic for matrix-matrix and matrix-vector product operations.

the storage requirements of these kernels and thus benefit from their high performance rates (as opposed to the performance of direct operation on sparse data structures). In particular, column three (column five) reports the ratio between the performance of SGEMM (SGEMV) and DGEMM (DGEMV). The BLAS libraries used are capable of exploiting the vector units where available, and thus the speedups shown in Table 4 are due to a combination of higher number of floating point operations completed at each clock cycle, reduced memory traffic on the bus, and higher cache hit rate.

Table 5 shows the difference in performance for the single and double precision implementation of the two sparse iterative solvers Conjugate Gradient and Generalized Minimal Residual. Columns two and three report the ratio between the performance of single and double precision CG for a fixed number (100) of iterations in both precondi-

Size	SCG/DCG		SGMRES/DGMRES	
	no prec.	prec.	no prec.	prec.
11,142	2.24	2.11	2.04	1.98
25,980	1.49	1.50	1.52	1.51
79,275	1.57	1.50	1.58	1.50
230,793	1.73	1.72	1.74	1.69
602,091	1.50	1.50	1.67	1.63

Table 5: Performance comparison between single and double precision arithmetic on a fixed number of iterations of Conjugate Gradient (100 iterations) and Generalized Minimal RESidual (2 cycles/restarts of GMRES(20)) methods both with and without Diagonal Scaling preconditioner. The runs were performed on Intel Woodcrest (3 GHz with a 1333 MHz front side bus).

Name	Reference BLAS			Goto BLAS		
	t_{SP}	t_{DP}	t_{DP}/t_{SP}	t_{SP}	t_{DP}	t_{DP}/t_{SP}
G64	102.76	80.54	1.27	101.27	85.23	1.18
Si10H16	392.71	308.83	1.27	343.88	345.18	0.99
c-71	1326.36	984.31	1.34	1266.28	1076.13	1.17
cage11	1274.90	945.81	1.34	1143.23	1047.29	1.09
dawson5	6.94	5.86	1.18	6.19	5.70	1.08
nasasrb	33.02	27.90	1.18	30.87	30.48	1.01
poisson3Db	553.19	404.26	1.36	515.00	450.44	1.14
rma10	3.10	2.74	1.13	2.97	2.52	1.17
wang4	21.77	16.99	1.28	19.32	18.96	1.01

Table 6: Time to solution of SuperLU in single and double precision for selected sparse matrices on Intel Woodcrest with reference and optimized BLAS.

tioned and unpreconditioned cases. Columns four and five report the same information for the GMRES(20) method where the number of cycles/restarts has been fixed to two. Since the sparse matrix kernels involved in these computations were not vectorized, the speedup shown in Table 5 is exclusively due to reduced data traffic on the bus and higher cache hit rate.

4.3 Experimental results

Figures 1- 4 show that the MUMPS single precision solver is always faster than the equivalent double precision one (i.e., the light bars are always above the thick horizontal line that corresponds to one). This is mainly due to both reduced data movement and better exploitation of vector arithmetic (via SSE2 or Altivec where present) since multifrontal methods have the ability to do matrix-matrix products.

The presented results also show that mixed precision iterative refinement is capable of providing considerable speedups with respect to the full double precision solver

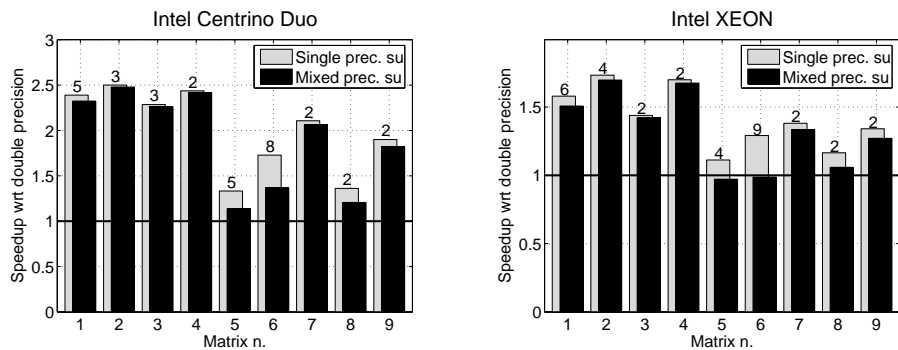


Figure 1: Experimental results for the MUMPS solvers. The light shaded bars reports the ratio between the performance of the single precision solver and the double precision one; the dark shaded bars report the ratio between the mixed precision solver and the double precision one. The number of iterations required to converge is given by the number above the bars. **Left:** Intel Centrino Duo. **Right:** Intel XEON.

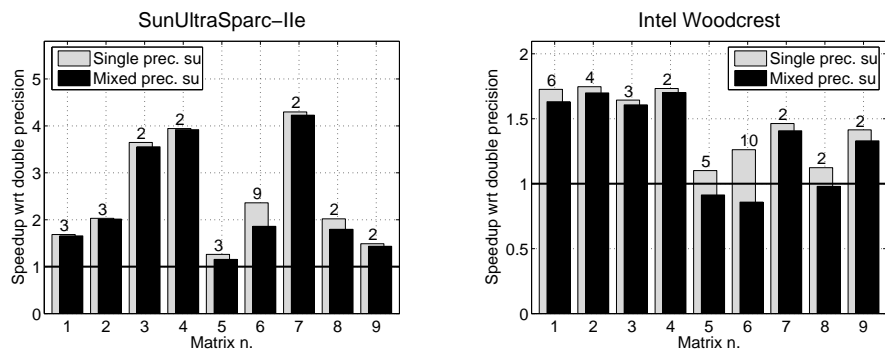


Figure 2: Experimental results for the MUMPS solvers. **Left:** Sun UltraSparc-IIe. **Right:** Intel Woodcrest.

while providing the same (in many cases also better) accuracy. To run these experiments, a convergence criterion different than that discussed in Section 3.1 was used. To make the comparison fair, in fact, the iterative refinement was stopped whenever the residual norm was the same as that computed for the full double precision solver. The performance of the mixed precision solver is usually very close to that of the single precision mainly because the cost of each iteration is often negligible if compared to the cost of the matrix factorization. It is important to note that, in some cases, the speedups reach very high values (more than 4.0 faster for the Poisson3Db matrix on the Sun UltraSparc-IIe architecture). This is due to the fact that the memory requirements are too high to accommodate the fill-in generated in the factorization phase as the main memory available on the system is lower for this machine. This forces the virtual memory system to swap pages to disk resulting in a considerable loss of performance;

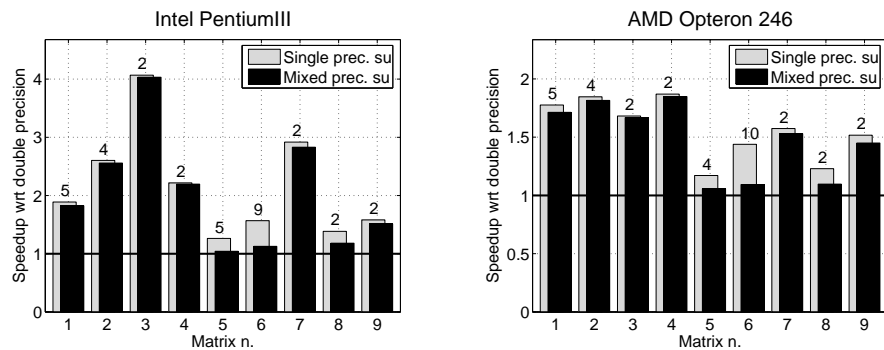


Figure 3: Experimental results for the MUMPS solvers. **Left:** Intel Pentium III Cop-
permine. **Right:** AMD Opteron 246.

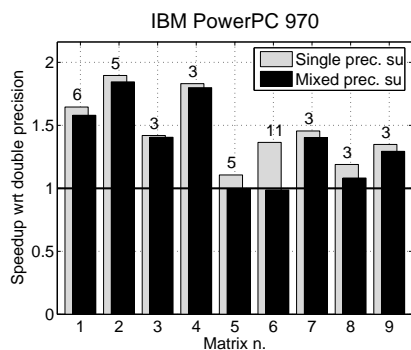


Figure 4: Experimental results for the MUMPS solvers. PowerPC 970.

since double precision data is twice as large as single precision, disk swapping may affect only the double precision solver and not the single precision one. It can be noted that disk swapping issues did not affect the results measured on those machines that are equipped with a bigger memory while it is usual on the Intel Pentium III and the Sun UltraSparc-IIe architectures that only have 512 MB of memory. Finally, the data presented in Figures 1- 4 show that, for some cases, the mixed precision iterative refinement solver does not provide a speedup. This can be mainly associated to three causes (or any combination of them):

1. The difference in performance between the single and the double precision solver is too small. In this case, even a few iterations of the refinement phase will compensate for the small speedup. This is, for example, the case of the Dawson5 matrix on the PowerPC 970 architecture.
2. The number of iterations to convergence is too high. The number of iterations to convergence is directly related to the matrix condition number [30]. The case of the Nasasrb matrix on the PowerPC 970 architecture show that even if the single

Matrix No.	Matrix name	DP residual	MP residual	# it.
1	G64	1.61e-10	2.00e-11	6
2	Si10H16	3.03e-12	1.93e-14	4
3	c-71	1.00e-14	9.78e-15	3
4	cage11	3.94e-16	1.04e-16	2
5	dawson5	5.88e-11	4.40e-12	5
6	nasasrb	1.08e-10	8.12e-11	10
7	poisson3Db	1.56e-14	1.04e-14	2
8	rma10	1.03e-13	6.85e-14	2
9	wang4	9.65e-15	6.13e-15	2

Table 7: Accuracy of the double precision solver and mixed precision, iterative solver on the Intel Woodcrest architecture. The last column reports the number of iterations performed by the mixed precision method to achieve the reported accuracy.

precision solver is almost $1.4\times$ faster, the mixed precision solver is slower than the double precision one because of the high number of iterations (11) needed to achieve the same accuracy. If the condition number is too high, the method may not converge at all.

3. The cost of each iteration is high compared to the performance difference between the double precision solver and the single precision one. In this case, even a few iterations can eliminate the benefits of performing the system factorization in single precision. As an example, take the case of the Rma10 matrix on the Intel Woodcrest architecture; two iteration steps on this matrix took 0.1 seconds, which is almost the same time needed to perform six iteration steps on the G64 matrix.

It is worth noting that, apart from the cases where the method does not converge, whenever the method results in a slowdown, the loss is on average only 7%.

Table 7 shows the residual of the solutions computed with the full double precision solver and the mixed precision, iterative solver for sparse direct methods on the Intel Woodcrest architecture. Note that for the matrices used and in general, for well-conditioned matrices, the mixed precision, iterative method is capable of delivering the same or better accuracy than the full double precision solver. The same was also observed for the mixed precision, sparse iterative solvers.

Table 6 shows timings of the sequential version of SuperLU on selected matrices from our test collection for single and double precision solvers. Both reference and Goto BLAS timings are shown. The sequential version of SuperLU implements matrix-vector multiply (`_GEMV`) as its computational kernel. This explains rather modest gains (if any) in the performance of the single precision solver over the double precision one: only up to 30%. The table also reveals that when optimized BLAS are used, the single precision is slower than double for some matrices, an artifact of small sizes of dense matrices passed to BLAS and the level of optimization of the BLAS for this particular architecture. The results are similar for other tested architectures, which

leads to a conclusion that there is not enough benefit in using our mixed precision approach for this version of SuperLU.

Next, we present our results on the mixed precision, iterative sparse solvers from Subsection 3.2. All the results are from runs on Intel Woodcrest (3 GHz with a 1333 MHz front side bus).

In Figure 5, we give the speedups for using mixed SP-DP vs DP-DP CG (dark bars). Namely, on the left we have the results for CG-PCG and on the right for PCG-PCG with diagonal preconditioner in the inner loop PCG. Similarly, in Figure 6, we give the results for GMRES-FGMRES (on the left) and PGMRES-FGMRES (on the right). Also, we compare the speedups of using SP vs DP for just the reference CG and PCG (correspondingly left and right in Figure 5; light bars), and SP vs DP for the reference GMRES and PGMRES (light bars in Figure 6). Note that in a sense the speedups in the light bars should represent the maximum that could be achieved by using the mixed precision algorithms. The fact that we get close to this maximum performance shows that we have successfully shifted the load from DP to SP arithmetic (with overall computation having less than 5% in DP arithmetic). The reason that the performance speedup for SP-DP vs DP-DP GMRES-FGMRES in Figure 6 (left, 4th matrix) is higher than the speedup of SP GMRES vs DP GMRES is that the SP-DP GMRES-FGMRES did one less outer cycle until convergence than the DP-DP GMRES-FGMRES.

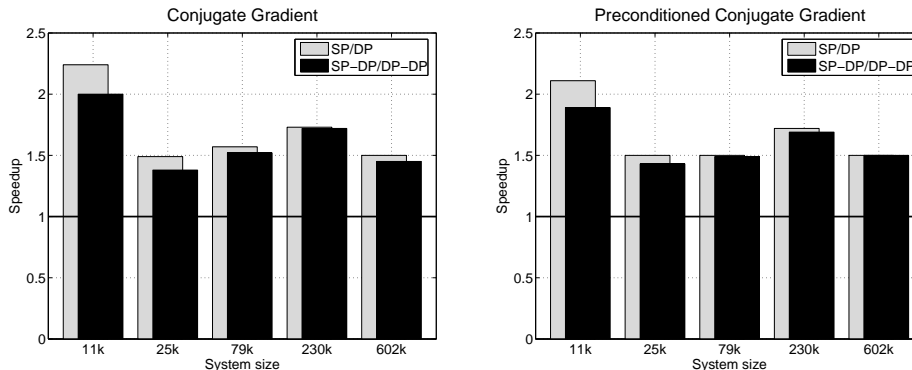


Figure 5: **Left:** Speedup of using SP vs DP CG (light bars) and the SP-DP vs DP-DP CG-PCG (dark). **Right:** Similar graph comparison but for the PCG algorithm (see also Subsection 3.2.1). The computations were on a Intel Woodcrest (3 GHz with a 1333 MHz front side bus).

Results comparing the SP-DP methods with the reference DP methods are shown in Figure 7 (left is a comparison for CG, right is for GMRES). The numbers on top of the bars on the left graph indicate the overhead, as the number of iterations, that took the mixed precision method to converge versus the reference DP method (e.g., overhead of 0.1 indicates 10% more iterations were performed in the mixed SP-DP vs the DP method). Even with the overhead, we see a performance speedup of at least 20% over the tested matrices. For the GMRES-based mixed precision methods we see

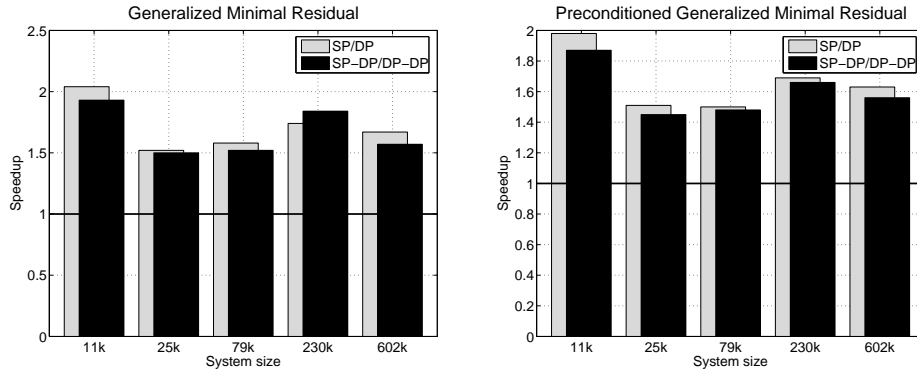


Figure 6: **Left:** Speedup of using SP vs DP GMRES (light bars) and the SP-DP vs DP-DP GMRES-FGMRES (dark). **Right:** Similar graph comparison but for the PGMRES algorithm (see also Subsection 3.2.2. The computations were on a Intel Woodcrest (3 GHz with a 1333 MHz front side bus).

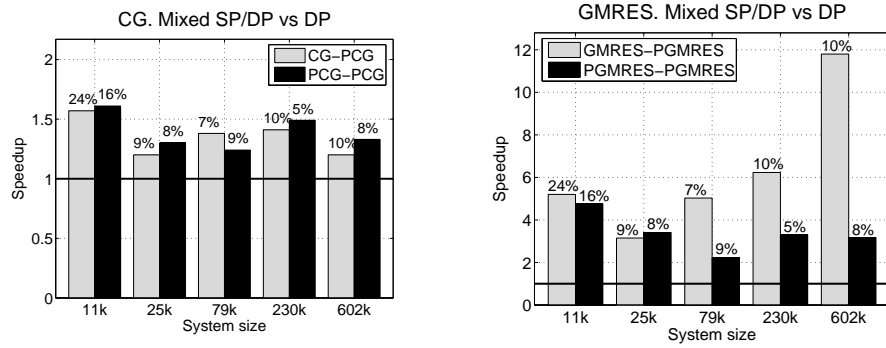


Figure 7: **Left:** Speedup of mixed SP-DP CG-PCG vs DP CG (light bars) and SP-DP PCG-PCD vs DP CG (dark) with diagonal preconditioner. **Right:** Similar graph comparison but for the GMRES based algorithms. The computations were on a Intel Woodcrest (3 GHz with a 1333 MHz bus).

a significant improvement, based on a reduced number of iterations and the effect of the SP speedup (from 45 to 100% as indicated in Figure 6). For example, the speedup factor of 12 for the biggest problem is due to speedup factors of approximately 7.5 from improved convergence and 1.6 from effects associated with the introduced SP storage and arithmetic.

It is not known how to choose the restart size m to get optimal results even for the reference GMRES (m). A reasonable working assumption is the bigger m the better, because one assumes that bigger m will get closer to the full GMRES. Following this assumption though does not guarantee better execution time, and sometimes the convergence can get even worse [20]. An alternative worth further exploration is to use a

truncated version of GMRES [39]. Another interesting approach is self adaptivity [16]. Here, to do a fair comparison, we ran it for $m = 25, 50$ (PETSc’s default [7]), 100, 150, 200, and 300, and chose the best execution time. Experiments show that the mixed precision method suggested is stable in regard to changing the restart values in the inner and outer loops. The experiments presented are for inner and outer $m = 20$. Note that this choice also results in less memory requirements than GMRES with $m \approx 70$ and higher (for most of the runs GMRES(100) and was best among the above choices for m), since the overhead in terms of DP vectors is $20 + 20$ (outer GMRES) $+ 10$ (20 SP vectors in the inner loop) $+ 20$ (matrix coefficients in SP; there are approximately 40 non-zeroes per row; see Table 2). In all the cases presented we had the number of inner cycles/restarts set to one.

Finally, we note the speedup for direct and iterative methods and its effect on performance. The speed up of moving to single precision for `_GEMM`-calling code (MUMPS) was approaching two and thus guaranteed success of our mixed precision, iterative refinement just as it did for the dense matrix operations. Not so for the `_GEMV`-calling code (SuperLU) for which the speedup did not exceed 30%, and thus no performance improvement was expected. However, for most of the iterative methods, the speedup was around 50%, and still we claim our approach to be successful. Inherently, the reason for speedup is the same for both settings (SuperLU and the iterative methods): the reduced memory bus traffic and possible super-linear effects when data fits in cache while being stored in single precision. But for the SuperLU case, there is the direct method overhead: the maintenance of evolving sparse data structures, which is done in fixed-point arithmetic so it does not benefit from using single precision floating-point arithmetic and hence yields the overall performance gains insufficient for our iterative refinement approach.

5 Future Work

We are considering a number of extensions and new directions for our work. The most broad category is the parallel setting. MUMPS is a parallel code, but it was used in a sequential setting in this study. Similarly, SuperLU has a parallel version that differs from the sequential counterpart in a very important way: it uses the matrix-matrix multiply kernel (`_GEMM`). This would give a better context for comparing multifrontal and supernodal approaches since they use the same underlying computational library. The only caveat is the lack of a single precision version of the parallel SuperLU solver. Another aspect brought by the latter solver is using static pivoting, while it vastly improves numerical stability of parallel SuperLU, it also improves the convergence of the iterative refinement that follows. This should result in less iterations and shorter solve time.

Using PETSc and its parallel framework for (among others) iterative methods could give us an opportunity to investigate our approach for a wider range of iterative methods and preconditioning scenarios. First though, we would have to overcome a technical obstacle of combining two versions of PETSc (one using single and one using double precision) in a single executable.

We have performed preliminary experiments on actual IBM Cell BE hardware (as

opposed to the simulator, which does not accurately account for memory system effects – a crucial component of sparse methods) with sparse matrix operations and are encouraged by the results to port our techniques in full. This would allow us to study their behavior with a much larger gap in the performance of the two precisions.

Our algorithms and their above descriptions focus solely on two precisions: single and double. We see them however in a broader context of higher and lower precision where, for example, a GPU performs computationally intensive operations in its native 16-bit arithmetic, and consequently the solution is refined using 128-bit arithmetic emulated in software (if necessary). As mentioned before, the limiting factor is conditioning of the system matrix. In fact, an estimate (up to the order of magnitude) of the condition number (often available from previous runs or the physical problem properties) may become an input parameter to an adaptive algorithm [17] that attempts to utilize the fastest hardware available, if its limited precision can guarantee convergence.

Also, the methods for sparse eigenvalue problems that result in Lanczos and Arnoldi algorithms are amenable to our techniques, and we would like to study their theoretical and practical challenges.

References

- [1] Patrick R. Amestoy, Iain S. Duff, and J.-Y. L'Excellent. Multifrontal parallel distributed symmetric and unsymmetric solvers. *Comput. Methods Appl. Mech. Eng.*, 184:501–520, 2000.
- [2] Patrick R. Amestoy, Iain S. Duff, J.-Y. L'Excellent, and Jacko Koster. A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM J. Matrix Anal. Appl.*, 23:15–41, 2001.
- [3] Patrick R. Amestoy, A. Guermouche, J.-Y. L'Excellent, and S. Pralet. Hybrid scheduling for the parallel solution of linear systems. *Parallel Comput.*, 32:136–156, 2006.
- [4] E. Anderson, Z. Bai, C. Bischof, Suzan L. Blackford, James W. Demmel, Jack J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and Danny C. Sorensen. *LAPACK User's Guide*. Society for Industrial and Applied Mathematics, Philadelphia, Third edition, 1999.
- [5] Cleve Ashcraft, R. Grimes, J. Lewis, Barry W. Peyton, and Horst Simon. Progress in sparse matrix methods in large sparse linear systems on vector supercomputers. *Intern. J. of Supercomputer Applications*, 1:10–30, 1987.
- [6] O. Axelsson and P. S. Vassilevski. A black box generalized conjugate gradient solver with inner iterations and variable-step preconditioning. *SIAM J. Matrix Anal. Appl.*, 12(4):625–644, 1991.
- [7] Satish Balay, Kris Buschelman, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, Lois Curfman McInnes, Barry F. Smith, and Hong Zhang. PETSc Web page, 2001. <http://www.mcs.anl.gov/petsc>.

- [8] R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. M. Donato, Jack Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. V. der Vorst. *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*. Philadelphia: Society for Industrial and Applied Mathematics., 1994. Also available as postscript file at <http://www.netlib.org/templates/Templates.html>.
- [9] Å. Björck. Iterative refinement and reliable computing. In M. G. Cox and S. Hammarling, editors, *Reliable Numerical Computation*, pages 249–266. Oxford University Press, Oxford, UK, 1990.
- [10] Alfredo Buttari, Jack Dongarra, Jakub Kurzak, Piotr Luszczek, and Stanmire Tomov. Computations to enhance the performance while achieving the 64-bit accuracy. Technical Report UT-CS-06-584, University of Tennessee Knoxville, November 2006. LAPACK Working Note 180.
- [11] Timothy A. Davis. A combined unifrontal/multifrontal method for unsymmetric sparse matrices. *ACM Transactions on Mathematical Software*, 25:1–19, 1999.
- [12] Timothy A. Davis. A column pre-ordering strategy for the unsymmetric-pattern multifrontal method. *ACM Transactions on Mathematical Software*, 30:196–199, 2004.
- [13] James W. Demmel. *Applied Numerical Linear Algebra*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1997.
- [14] James W. Demmel, Stanley C. Eisenstat, John R. Gilbert, Xiaoye S. Li, and Joseph W. H. Liu. A supernodal approach to sparse partial pivoting. *SIAM J. Matrix Anal. Appl.*, 20(3):720–755, 1999.
- [15] James W. Demmel, John R. Gilbert, and Xiaoye S. Li. A asynchronous parallel supernodal algorithm for sparse gaussian elimination. *SIAM J. Matrix Anal. Appl.*, 20(3):915–952, 1999.
- [16] Jim Demmel, Jack Dongarra, Victor Eijkhout, Erika Fuentes, Antoine Petit, Rich Vuduc, R. Clint Whaley, and Kathy Yelick. Self-adapting linear algebra algorithms and software. *Proceedings of the IEEE*, 93(2), February 2005. See <http://www.spiral.net/ieee-special-issue/overview.html>.
- [17] Jack J. Dongarra and Victor Eijkhout. Self-adapting numerical software for next generation applications. Technical Report Lapack Working Note 157, ICL-UT-02-07, Innovative Computing Lab, University of Tennessee, August 2002. To appear IJHPCA 17(2) 2003, <http://icl.cs.utk.edu/iclprojects/pages/sans.html>.
- [18] Iain S. Duff and John K. Reid. The multifrontal solution of indefinite sparse symmetric linear equations. *ACM Transactions on Mathematical Software*, 9(3):302–325, September 1983.

- [19] Timothy A. Davis Iain S. Duff. An unsymmetric-pattern multifrontal method for sparse LU factorization. *SIAM J. Matrix Anal. Appl.*, 18:140–158, 1997.
- [20] Mark Embree. The tortoise and the hare restart gmres. *SIAM Review*, 45:259–266, 2003.
- [21] George E. Forsythe and Cleve B. Moler. *Computer Solution of Linear Algebraic Systems*. Prentice-Hall, Englewood Cliffs, NJ 07632, USA, 1967.
- [22] D. Göddeke, R. Strzodka, and S. Turek. Accelerating double precision FEM simulations with GPUs. In F. Hülsemann, M. Kowarschik, and U. Rüde, editors, *Simulationstechnique 18th Symposium in Erlangen, September 2005*, volume Frontiers in Simulation, pages 139–144. SCS Publishing House e.V., 2005. ASIM 2005.
- [23] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, USA, second edition, 1989.
- [24] Gene H. Golub and Qiang Ye. Inexact preconditioned conjugate gradient method with inner-outer iteration. *SIAM Journal on Scientific Computing*, 21(4):1305–1320, 2000.
- [25] W. D. Gropp, D. K. Kaushik, D. E. Keyes, and B. F. Smith. Latency, bandwidth, and concurrent issue limitations in high-performance cfd. Technical Report ANL/MCS-P850-1000, Argonne National Laboratory, October 2000.
- [26] William D. Gropp, Dinesh K. Kaushik, David E. Keyes, and Barry F. Smith. High-performance parallel implicit CFD. *Parallel Computing*, 27(4):337–362, 2001.
- [27] W. Hackbusch. *Multigrid Methods and Applications*. Springer Series in Computational Mathematics Vol. 4, Springer-Verlag, Berlin, 1985.
- [28] Nicholas J. Higham. *Accuracy and Stability of Numerical Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, second edition, 2002.
- [29] Nicholas J. Higham. *Accuracy and Stability of Numerical Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, second edition, August 2002.
- [30] Julie Langou, Julien Langou, Piotr Luszczek, Jakub Kurzak, Alfredo Buttari, and Jack Dongarra. Exploiting the performance of 32 bit floating point arithmetic in obtaining 64 bit accuracy. In *Proceedings of SC06*, Tampa, Florida, November 11-17 2006. See <http://icl.cs.utk.edu/iter-ref>.
- [31] Xiaoye S. Li. *Sparse Gaussian Elimination on High Performance Computers*. Computer Science Department, University of California at Berkeley, 1996. Ph.D. thesis (SuperLU software available at <http://www.nersc.gov/~xiaoye/SuperLU/>).

- [32] Xiaoye S. Li and James W. Demmel. SuperLU_DIST: A scalable distributed-memory sparse direct solver for unsymmetric linear systems. *ACM Transactions on Mathematical Software*, 29:110–140, 2003.
- [33] C. B. Moler. Iterative refinement in floating point. *Journal of the Association for Computing Machinery*, 14(2):316–321, April 1967.
- [34] Y. Notay. Flexible conjugate gradients. *SIAM Journal on Scientific Computing*, 22:1444–1460, 2000.
- [35] A. Quarteroni and A. Valli. *Domain Decomposition Methods for Partial Differential Equations*. Oxford University Press, 1999.
- [36] Y. Saad. A flexible inner-outer preconditioned GMRES algorithm. Technical Report 91-279, Department of Computer Science and Engineering, University of Minnesota, Minneapolis, Minnesota, 1991.
- [37] Y. Saad. *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2003.
- [38] Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual method for solving nonsymmetric linear systems. *SIAM J. Sci. Statist. Comput.*, pages 856–869, 1986.
- [39] Yousef Saad and Kesheng Wu. DQGMRES: a direct quasi-minimal residual algorithm based on incomplete orthogonalization. *Numerical linear algebra with applications*, 3(4):329–343, 1996.
- [40] V. Simoncini and D. Szyld. Theory of inexact krylov subspace methods and applications to scientific computing, 2002.
- [41] Valeria Simoncini and Daniel B. Szyld. Flexible inner-outer Krylov subspace methods. *SIAM J. Numer. Anal.*, 40(6):2219–2239, 2002.
- [42] Valeria Simoncini and Daniel B. Szyld. The effect of non-optimal bases on the convergence of krylov subspace methods. *Numer. Math.*, 100(4):711–733, 2005.
- [43] G. W. Stewart. *Matrix algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2001.
- [44] R. Strzodka and D. Goddeke. Mixed precision methods for convergent iterative schemes, May 2006. EDGE 2006, 23.-24. May 2006, Chapel Hill, North Carolina.
- [45] R. Strzodka and D. Goddeke. Pipelined mixed precision algorithms on FPGAs for fast and accurate PDE solvers from low precision components. In *IEEE Proceedings on Field-Programmable Custom Computing Machines (FCCM 2006)*. IEEE Computer Society Press, May 2006. to appear.
- [46] Kathryn Turner and Homer F. Walker. Efficient high accuracy solutions with gmres(m). *SIAM J. Sci. Stat. Comput.*, 13(3):815–825, 1992.

- [47] University of Florida sparse matrix collection.
<http://www.cise.ufl.edu/research/sparse/matrices/>.
- [48] J. van den Eshof, G. L. G. Sleijpen, and M .B. van Gijzen. Relaxation strategies for nested Krylov methods. Technical Report TR/PA/03/27, CERFACS, Toulouse, France, 2003.
- [49] H. A. van der Vorst and C. Vuik. GMRESR: a family of nested GMRES methods. *Numerical Linear Algebra with Applications*, 1(4):369–386, 1994.
- [50] C. Vuik. New insights in gmres-like methods with variable preconditioners. *J. Comput. Appl. Math.*, 61(2):189–204, 1995.
- [51] J. H. Wilkinson. *The Algebraic Eigenvalue Problem*. Oxford University Press, Oxford, UK, 1965.