

Accelerating Restarted GMRES With Mixed Precision Arithmetic

Neil Lindquist¹, Piotr Luszczek², and Jack Dongarra³, *Fellow, IEEE*

Abstract—The generalized minimum residual method (GMRES) is a commonly used iterative Krylov solver for sparse, non-symmetric systems of linear equations. Like other iterative solvers, data movement dominates its run time. To improve this performance, we propose running GMRES in reduced precision with key operations remaining in full precision. Additionally, we provide theoretical results linking the convergence of finite precision GMRES with classical Gram-Schmidt with reorthogonalization (CGSR) and its infinite precision counterpart which helps justify the convergence of this method to double-precision accuracy. We tested the mixed-precision approach with a variety of matrices and preconditioners on a GPU-accelerated node. Excluding the incomplete LU factorization without fill in (ILU(0)) preconditioner, we achieved average speedups ranging from 8 to 61 percent relative to comparable double-precision implementations, with the simpler preconditioners achieving the higher speedups.

Index Terms—Linear systems, multiple precision arithmetic

1 INTRODUCTION

THE generalized minimum residual method (GMRES) is an iterative solver for sparse, non-symmetric systems of linear equations [1]. Like most iterative solvers, GMRES consists mostly of matrix-vector and vector-vector operators, resulting in a low computational intensity. Hence, reducing data movement is necessary to improve performance. Towards this end, we present an implementation of GMRES using a mix of single- and double-precision that is designed to achieve the same final accuracy as double-precision GMRES while reducing data movement.

In short, GMRES constructs a basis for the Krylov subspace using Arnoldi's procedure [2], then finds the solution vector from that subspace that minimizes the 2-norm of the resulting residual. One particularly useful modification to GMRES is restarting, which is used to limit the memory usage and computational requirements of the growing Krylov basis [3]. We focus on two schemes for orthogonalizing the Krylov basis in the Arnoldi procedure, modified Gram-Schmidt (MGS) and classical Gram-Schmidt with reorthogonalization (CGSR). The former is often used due to its lower computational and data-access costs [4], while the latter better retains orthogonality and can be implemented using matrix-vector products [5]. For CGSR, we always reorthogonalize once, which is sufficient to provide

numerical stability [6]. Additionally, we focused on a left-preconditioned version of GMRES but expect the results to hold for right-preconditioned versions too. Algorithm 1 shows the formulation of GMRES we used.

Based on our previous work, we focus on a specific mixed-precision strategy that has been successful regarding both accuracy and CPU performance [7]. This approach uses single precision everywhere except to compute the residual and update the solution. Specifically, it uses double precision for lines 3 and 25 in Algorithm 1 and uses single precision for everything else. (I.e., computation done in double precision is labeled with u_{high} and computation done in single precision is labeled with u_{low}). This requires a copy of A in single precision, and two vector type conversions per restart (z_k and u_k) but has the advantage of otherwise using existing uniform-precision kernels. Our previous work included experiments with the effects on convergence when reducing various parts of GMRES and using different restart strategies; those experiments guided the design of our mixed-precision scheme. Herein, we extend the support for this approach by providing stronger theoretical analysis of the reduced precision inner solver and new experimental results on a GPU-accelerated node with a variety of preconditioners.

2 PREVIOUS WORK

The idea of using multiple precisions to improve performance has been applied successfully to a variety of problems, particularly in recent years [8]. Furthermore, it is a well-established method for improving the performance of solving systems of linear equations, particularly for dense linear systems [9], [10].

One approach to using multiple precisions in GMRES is to store the preconditioner in reduced precision and doing the rest of the computation in full precision [11]. The approximate nature of the preconditioner means that reducing the floating-point precision has a limited reduction in

- Neil Lindquist and Piotr Luszczek are with Innovative Computing Laboratory, University of Tennessee, Knoxville, TN 37996 USA. E-mail: {nlindqu1, luszczek}@icl.utk.edu.
- Jack Dongarra is with Innovative Computing Laboratory, University of Tennessee, Knoxville, TN 37996 USA, and with Oak Ridge National Laboratory, Oak Ridge, TN 37831 USA, and also with the University of Manchester, M13 9PL Manchester, U.K. E-mail: dongarra@icl.utk.edu.

Manuscript received 24 Feb. 2021; revised 15 June 2021; accepted 16 June 2021.
Date of publication 22 June 2021; date of current version 15 Oct. 2021.

(Corresponding author: Neil Lindquist.)

Recommended for acceptance by S. Alam, L. Curfman McInnes, and K. Nakajima.
Digital Object Identifier no. 10.1109/TPDS.2021.3090757

quality. One interesting variation of this is to precondition a full-precision GMRES with a reduced-precision GMRES (possibly with a preconditioner of its own) [12]. This is similar to our approach, but we use iterative refinement as the outer solver instead of GMRES.

Algorithm 1. Restarted GMRES in Mixed-Precision With Left Preconditioning [3]

```

1:  $A \in \mathbb{R}^{n \times n}$ ,  $x_0, b \in \mathbb{R}^n$ ,  $M^{-1} \approx A^{-1}$ 
2: for  $k = 1, 2, \dots$  do
3:    $z_k \leftarrow b - Ax_k$  ( $u_{\text{high}}$ )
4:   If  $\|z_k\|_2$  is small enough, stop ( $u_{\text{high}}$ )
5:    $r_k \leftarrow M^{-1}z_k$  ( $u_{\text{low}}$ )
6:    $\beta \leftarrow \|r_k\|_2$ ,  $s_0 \leftarrow \beta$  ( $u_{\text{low}}$ )
7:    $v_1 \leftarrow r_k/\beta$ ,  $V_1 \leftarrow [v_1]$  ( $u_{\text{low}}$ )
8:    $j \leftarrow 0$ 
9:   loop until the restart condition is met
10:   $j \leftarrow j + 1$ 
11:   $w \leftarrow M^{-1}Av_j$  ( $u_{\text{low}}$ )
12:   $w, h_{1,j}, \dots, h_{j,j} \leftarrow \text{orth}(w, V_j)$   $\triangleright$  MGS or CGSR
13:   $h_{j+1,j} \leftarrow \|w\|_2$  ( $u_{\text{low}}$ )
14:   $v_{j+1} \leftarrow w/h_{j+1,j}$  ( $u_{\text{low}}$ )
15:   $V_{j+1} \leftarrow [V_j, v_{j+1}]$  ( $u_{\text{low}}$ )
16:  for  $i = 1, \dots, j - 1$  do
17:     $\begin{bmatrix} h_{i,j} \\ h_{i+1,j} \end{bmatrix} \leftarrow \begin{bmatrix} \alpha_i & \beta_i \\ -\beta_i & \alpha_i \end{bmatrix} \times \begin{bmatrix} h_{i,j} \\ h_{i+1,j} \end{bmatrix}$  ( $u_{\text{low}}$ )
18:  end for
19:   $\begin{bmatrix} \alpha_j \\ \beta_j \end{bmatrix} \leftarrow \text{rotation\_matrix}\left(\begin{bmatrix} h_{j,j} \\ h_{j+1,j} \end{bmatrix}\right)$  ( $u_{\text{low}}$ )
20:   $\begin{bmatrix} s_j \\ s_{j+1} \end{bmatrix} \leftarrow \begin{bmatrix} \alpha_j & \beta_j \\ -\beta_j & \alpha_j \end{bmatrix} \times \begin{bmatrix} s_j \\ 0 \end{bmatrix}$  ( $u_{\text{low}}$ )
21:   $\begin{bmatrix} h_{j,j} \\ h_{j+1,j} \end{bmatrix} \leftarrow \begin{bmatrix} \alpha_j & \beta_j \\ -\beta_j & \alpha_j \end{bmatrix} \times \begin{bmatrix} h_{j,j} \\ h_{j+1,j} \end{bmatrix}$  ( $u_{\text{low}}$ )
22:  end loop
23:   $H \leftarrow \{h_{i,\ell}\}_{1 \leq i, \ell \leq j}$ ,  $s \leftarrow [s_1, \dots, s_j]^T$ 
24:   $u_k \leftarrow V_j H^{-1}s$  ( $u_{\text{low}}$ )
25:   $x_{k+1} \leftarrow x_k + u_k$  ( $u_{\text{high}}$ )
26: end for
27: procedure MGS $w, V_j$ 
28:    $[v_1, \dots, v_j] \leftarrow V_j$ 
29:   for  $i = 1, 2, \dots, j$  do
30:      $h_{i,j} \leftarrow w \cdot v_i$  ( $u_{\text{low}}$ )
31:      $w \leftarrow w - h_{i,j}v_i$  ( $u_{\text{low}}$ )
32:   end for
33:   return  $w, h_{1,j}, \dots, h_{j,j}$ 
34: end procedure
35: procedure CGSR $w, V_j$ 
36:    $h \leftarrow V_j^T w$  ( $u_{\text{low}}$ )
37:    $w \leftarrow w - V_j h$  ( $u_{\text{low}}$ )
38:    $g \leftarrow V_j^T w$  ( $u_{\text{low}}$ )
39:    $w \leftarrow w - V_j g$  ( $u_{\text{low}}$ )
40:    $[h_{0,j}, \dots, h_{j,j}]^T \leftarrow h + g$  ( $u_{\text{low}}$ )
41:   return  $w, h_{1,j}, \dots, h_{j,j}$ 
42: end procedure

```

There has recently been useful theoretical work for varying the working precision as the iteration progresses in a non-restarted GMRES [13]. Notably, part of this work shows

that computing the Arnoldi process in finite-precision allows GMRES to converge at approximately the same rate as GMRES computed exactly until an accuracy related to round-off error is reached. Our theoretical results in Section 3 are based on some of these ideas. There exists further theoretical work on reducing the accuracy of just the matrix vector products in GMRES and other Krylov solvers as the number of inner iterations progresses [14], [15], [16]. These approaches may avoid the need to restart, unlike our work; however, they increase the complexity of implementation and achieving high accuracy may require estimating the smallest singular value.

For restarted GMRES, there have been a few works involving using multiple precisions in the GMRES algorithm. The first is using mixed-precision iterative refinement where the inner solver is a single-precision GMRES [17], [18]. This approach is similar to what we tested; however, that work tests only limited configurations of GMRES and matrices. The second approach is to store just the Krylov basis in reduced precision, which was tested with both floating- and fixed-point formats and 32- and 16-bits per value [19]. It was successful in providing a speedup with the 32-bit floating-point version providing the best median speedup at 1.4 times. However, the scheme, as described, requires custom, high-performance, mixed-precision kernels, which increases the cost of implementation due to the limited availability of existing mixed-precision routines.

One final work with GMRES is the use of integer arithmetic [20]. While integer GMRES did not involve a reduction in data movement, it does show GMRES achieving a full-precision solution with limited iteration overhead when the solver uses an alternative data format. Relatedly, there has been work to use data compression techniques in flexible GMRES, although only for the non-orthogonalized Krylov vectors [21].

Mixed precision approaches have also been used for other iterative solvers. Like GMRES, mixed precision approaches include a reduced precision preconditioner [22], [23] and using a reduced precision solver inside iterative refinement [17]. However, with Krylov methods, iterative refinement discards the subspace at each restart; so, the strategy of “reliable updates” has been proposed, which retains information about the Krylov subspace across restarts [24], [25]. Finally, there has been some exploration of the use of alternate data representations, such as data compression, in iterative solvers [26], [27], [28].

3 NUMERICAL PROPERTIES OF MIXED PRECISION GMRES

It is important to understand the effects of reductions in precision on the accuracy of the final solutions. Note that restarted GMRES is equivalent to iterative refinement using a non-restarted GMRES as the inner solver, which provides a powerful tool for recovering accuracy. This equivalence can be seen by noting that lines 5 through 24 in Algorithm 1 are equivalent to a non-restarted GMRES with an initial guess of x_k and the remaining form iterative refinement of the non-restarted GMRES.

First, consider the parts of the solver that must remain in full precision. Notably, the linear system being solved, $Ax = b$, must be stored in full precision. In general, reductions in the accuracy of A and b directly change the problem being solved and can introduce a backward error on the order of the reduced precision. If x is stored in reduced precision, a similar forward error will be introduced. As an extension of this, adding the updates from the inner solve to x must be done in full precision, to ensure x remains in full precision. Finally, the residual $r = b - Ax$ must be computed in full precision, otherwise errors smaller than the reduced precision accuracy cannot be corrected.

Next, consider the effects of reducing precision in the computation of the error correction. For stationary iterative refinement, it is well-known that the error correction can be computed in reduced precision while still achieving a full precision final solution [9]. For GMRES, the theory is less developed. Note that single-precision GMRES is backwards stable to single precision [29], [30]. Then, mixed-precision iterative refinement can use this GMRES to compute a double-precision backwards-stable solution [31]. Thus, the approach should be backwards-stable to full-precision. However, this analysis ignores the possibility of restarting before a single-precision accurate solution is achieved which is a significant issue in practice due to the increasing memory cost of GMRES.

Recent work has shown that MGS-GMRES converges at approximately the same rate whether the Arnoldi process is computed in finite precision or exactly until the relative residual is below a roundoff-dependent threshold [13]. Unfortunately, it is difficult to turn this into a useful general-purpose bound. Towards this end, we provide the following theorem for CGSR-GMRES which takes advantage of the better orthogonality of CGSR to describe the accuracy of the finite precision solution relative to the exact precision one [5].

Theorem 1. Let \bar{x}_j and $x_j^{(e)}$ be the solutions computed by j iterations of finite precision GMRES and exact GMRES, respectively, without restarting. Let b be the right-hand side vector. Suppose $u < 10^{-3}$ and $c_4(n, j)u\kappa(A\bar{V}_j) < 1$ for a particular $c_4(n, j) \in \mathcal{O}(n^2 j^3)$ with \bar{V}_j being the computed Krylov basis. Let $\delta_+ = (1 + \sqrt{u})^{1/2}$ and $\delta_- = (1 - \sqrt{u})^{1/2}$. Then

$$\begin{aligned} \|b - A\bar{x}_j\|_2 &\leq \delta_+^2 \delta_-^2 \|b - Ax_j^{(e)}\|_2 \\ &\quad + 9\delta_+ u j \|A\|_2 \|\bar{x}_j\|_2 \\ &\quad + \delta_+^2 \delta_-^{-1} \gamma_p j^{1/2} \|A\|_F \|x_j^{(w)}\|_2 \\ &\quad + \delta_+ \delta_-^{-1} c_1(n, j) u \|A\|_F \|x_j^{(w)}\|_2 \\ &\quad + \frac{c_1(n, j) u \|A\|_2 \|\bar{x}_j\|_2}{\delta_- - \gamma_j j^{1/2} \delta_+} \\ &\quad + \frac{j^{1/2} \delta_+ \|A\|_F \|\bar{x}_j\|_2}{\delta_- - \gamma_j j^{1/2} \delta_+} \\ &\quad \times \left(\gamma_p + \gamma_j + (\gamma_j + 9u j \gamma_j + 9u j^{1/2}) \right. \\ &\quad \left. \times \delta_+ \delta_-^{-1} (2 + \gamma_p) \right). \end{aligned}$$

where $c_1(n, j) \in \mathcal{O}(nj)$ and $x_j^{(w)}$ is the solution computed by j iterations of a particular weighted-GMRES.

The polynomials $c_1(n, j)$ and $c_4(n, j)$ are the same as those in Giraud *et al.*'s error analysis of CGSR [6]. Note that \bar{x}_j and $x_j^{(e)}$ are equivalent to u_k from Algorithm 1 when the restart condition is met after j inner iterations.

When $u = 2^{-24}$, $\|x_j^{(e)}\| \approx \|\bar{x}_j\| \approx \|x_j^{(w)}\|$, $j^{3/2}u \ll 1$, and $p^{3/2}u \ll 1$, this can be simplified to

$$\begin{aligned} &\frac{\|b - A\bar{x}_j\|_2}{\|A\|_F \|\bar{x}_j\|_2 + \|b\|_2} \\ &\lesssim 1.1 \frac{\|b - Ax_j^{(e)}\|_2}{\|A\|_F \|x_j^{(e)}\|_2 + \|b\|_2} \\ &\quad + (4j^{3/2} + 30j + 3pj^{1/2} + 3c_1(n, j))u. \end{aligned}$$

In other words, finite precision CGSR-GMRES converges at effectively the same rate as its exact counterpart until an error of approximately $(4j^{3/2} + 30j + 3pj^{1/2} + 3c_1(n, j))u$ is reached. This implies that for restarted GMRES, if full precision GMRES can converge, then reduced precision GMRES should either converge similarly or satisfy the backward error threshold. In the latter case, the backward stability of iterative refinement will result in a backward stable solution [31]. In the former cases, we expect similar behavior, but differences in vector directions may reduce effectiveness when the solver restarts.

It should be noted that the low-order polynomials c_1 and c_4 can quickly become onerous for realistically sized matrices. Fortunately, the bounds provided by this theorem are much worse than will occur in practice. First, they assume round-off error will always accumulate without cancellation. However, the recent work on probabilistic error bounds for dot-products has shown that a relative error of about $u\sqrt{n}$ holds with probability close to 1, compared the formal worst-case bound of un [32]. Second, the error bounds assume that the dot product summation is done sequentially. However, GPU accelerated systems distribute the work across many threads which helps reduce the accumulation of errors [33].

4 IMPLEMENTATION CONSIDERATIONS

Periodic restarting is a key component in obtaining full precision accuracy in this approach as it relies on iterative refinement. For many problems, both the memory constraints and the increasing computational load will force the solver to restart before the accuracy of the inner solver is achieved. However, for some problems, GMRES can produce a full precision solution in relatively few iterations; in these cases, waiting to restart until after a fixed number of inner iterations will result in a stalled improvement when the round-off error has overwhelmed any meaningful contributions to the computed solution. We have previously discussed and tested various restart strategies [7]. In that work, we found that an effective restart strategy was to initiate the first restart when the residual approximation improves by a factor of 10^{-6} , then initiate subsequent restarts after the same number of inner iterations.

Total memory usage is an important constraint, particularly when considering the smaller memory sizes of GPU

TABLE 1
Properties of the Matrices Tested Without a Preconditioner

Matrix	Rows	Nonzeros	Condition	RHS
			Lower Bound	Provided
af_0_k101	5.0×10^5	1.8×10^7	5.5×10^5 †	yes
af_shell9	5.0×10^5	1.8×10^7	1.2×10^6 †	yes
apache2	7.2×10^5	4.8×10^6	3.0×10^6 †	no
atmosmodj	1.3×10^6	8.8×10^6	6.4×10^3	yes
BenElechi1	2.5×10^5	1.3×10^7	1.3×10^6 †	yes
bone010	9.9×10^5	4.8×10^7	1.6×10^6 †	no
Bump_2911	2.9×10^6	1.3×10^8	7.5×10^6 †	no
cage13	4.5×10^5	7.5×10^6	1.1×10^1	no
cage14	1.5×10^6	2.7×10^7	9.6×10^0	no
crankseg_1	5.3×10^4	1.1×10^7	1.4×10^7 †	no
CurlCurl_2	8.1×10^5	8.9×10^6	4.1×10^5 †	no
CurlCurl_4	2.4×10^6	2.7×10^7	3.4×10^5 †	no
ecology2	1.0×10^6	5.0×10^6	3.2×10^7 †	no
F1	3.4×10^5	2.7×10^7	7.1×10^5 †	yes
FEM_3D_thermal2	1.5×10^5	3.5×10^6	2.5×10^3	no
G3_circuit	1.6×10^6	7.7×10^6	6.0×10^6 †	no
hood	2.2×10^5	9.9×10^6	3.8×10^5 †	no
language	4.0×10^5	1.2×10^6	5.9×10^2	no
marinel	4.0×10^5	6.2×10^6	3.8×10^5 †	yes
mc2depi	5.3×10^5	2.1×10^6	1.3×10^{14}	no
ns3Da	2.0×10^4	1.7×10^6	5.6×10^2	yes
parabolic_fem	5.3×10^5	3.7×10^6	2.1×10^5 †	yes
poisson3Db	8.6×10^4	2.4×10^6	2.6×10^3	yes
pwtk	2.2×10^5	1.2×10^7	6.9×10^5 †	no
rajat31	4.7×10^6	2.0×10^7	4.0×10^6	no
stomach	2.1×10^5	3.0×10^6	2.9×10^1	no
t2em	9.2×10^5	4.6×10^6	2.2×10^5 †	no
thermal2	1.2×10^6	8.6×10^6	1.5×10^6 †	yes
tmt_unsym	9.2×10^5	4.6×10^6	2.3×10^8 †	no
torso2	1.2×10^5	1.0×10^6	2.0×10^1	no
torso3	2.6×10^5	4.4×10^6	9.5×10^1	no
venkat01	6.2×10^4	1.7×10^6	1.3×10^5 †	yes

†Condition estimator reached 200 000 iterations before the LSQR convergence criteria.

accelerators. Consider GMRES implemented for matrices in compressed sparse row (CSR) format and restarting after, at most, m inner iterations. The matrix entries, right-hand side, and solution take combined $12n_{nz} + \mathcal{O}(n)$ bytes. The double-precision GMRES requires an additional $8nm + \mathcal{O}(n + m^2)$ bytes. Instead, the mixed precision variant requires an additional $4n_{nz} + 4nm + \mathcal{O}(n + m^2)$ bytes. For problems with many nonzeros per row relative to m , the mixed-precision approach will require a larger total allocation, which may be onerous on memory-constrained systems. Storing the low-order and high-order bytes separately would circumvent this issue [34]. Regardless of matrix storage, the mixed-precision approach always has a smaller increase in allocation for an increase in m . On some large problems, this may allow for a larger basis and, thus, take better advantage of superlinear convergence in GMRES [35].

5 EXPERIMENTAL RESULTS

To test the performance of our approach on GPUs, we implemented a restarted GMRES using the Kokkos performance portability library [36] and NVIDIA's cuBLAS and cuSPARSE libraries. Kokkos was chosen for ease of use and is not expected to perform significantly different

from a CUDA implementation. To limit expensive memory transfers between CPU and GPU, all computation is done on the GPU and only the high level control flow is done on the CPU.

Various matrices in CSR format from the SuiteSparse collection with more than a million nonzero elements were tested [37]. Furthermore, we only used matrices that converged with fewer than 300 restarts for a given configuration. If a file ending with `_b` was provided by SuiteSparse, the first column was used as the right-hand side. For other matrices, the right-hand side was computed from a solution where each element was randomly selected from the uniform range $[0, 1)$. The matrices tested are shown in Table 1. In addition to structural properties, the table contains lower bounds of the condition numbers of these matrices, computed by testing forward error vectors of LSQR [38]. Many of the poorly conditioned matrices reached the iteration limit before LSQR's convergence criteria, so they may have significantly worse conditioning than these lower bounds imply.

We tested the effectiveness of the mixed-precision approach for a variety of preconditioners. First is the identity matrix to test unpreconditioned GMRES. While in practice some form of preconditioner is almost always used, not using a preconditioner makes it easier to compute




Fig. 1. Relative performance of unpreconditioned GMRES with MGS (top) or CGSR (bottom).

condition numbers. Second is a scalar Jacobi preconditioner. This is a diagonal matrix where each element is the inverse of the corresponding diagonal element of A to provide a measure of row scaling. Third is an ILU(0). This is a powerful preconditioner formed through Gaussian Elimination, except only entries that are nonzero in A are computed. However, the sparse triangular solves needed to apply the factorization cannot be parallelized to efficiently use a GPU. So, fourth, we test ILU(0), except with the triangular solves replaced with Jacobi iterations [39], [40]. Because we are focusing on how well the mixed-precision approach works for different types of preconditioners, we test each preconditioner in isolation and do not compare them. Furthermore, each preconditioner is computed in double precision, then converted to the appropriate precision to limit differences caused by preconditioner differences. In addition to testing double-precision GMRES and the proposed mixed-precision GMRES, we also tested a completely single-precision implementation and reducing the precision of just the preconditioner. In order to use existing uniform-precision kernels, the input and output of the preconditioner are converted to single precision.

Each test was run to reach a backward error of

$$\frac{\|b - Ax\|_2}{\|A\|_F \|x\|_2 + \|b\|_2} \leq 10^{-10}.$$

Any test requiring more than 300 restarts was considered a failure, with the tested matrices chosen to all succeed using the double-precision implementation. At most 100 inner

iterations were run before restarting, with three restart strategies tested:

- 1) just the inner-iteration count,
- 2) the residual approximation improving by a factor of 10^{-10} , and
- 3) the residual approximation improving by a factor of 10^{-6} for the first restart and the same number of inner iterations after that.

Note that mixed precision was only tested with the third strategy to prevent the choice of restart strategy from inappropriately benefiting it. For each configuration we ran the code three times, plotting the median, with error bars for the minimum and maximum. Run times include constructing the preconditioner and any type conversions. The restart strategy with the smallest median was plotted. Speedups were computed as the inverse of the geometric mean of the normalized mixed-precision times.

Our implementation is available at bitbucket.org/icl/mixed-precision-gmres under the tag TPDS. We used Kokkos 3.1.01, CUDA 10.2.89, MKL 2019.3.199, and GCC 7.3.0. Code was run on a machine with a single NVIDIA V100 GPU and two Haswell 10-core Intel(R) Xeon(R) CPU E5-2650 v3 @ 2.30GHz processors. Each CPU core has a 32 KiB L1 instruction cache, a 32 KiB L1 data cache, and a 256 KiB L2 cache. Each processor has a single 25 MiB L3 cache, and the entire node has 32 GiB of memory. The V100 card has 80 streaming multiprocessors, a 128 KiB L1 cache for each multiprocessor, a 6 MiB shared L2 cache, and a 16 GiB memory.

TABLE 2
Inner Iteration Counts for Unpreconditioned GMRES

Matrix	Double		Mixed		Single	
	MGS	CGSR	MGS	CGSR	MGS	CGSR
apache2	21 400	21 400	21 500	21 500	29 200	29 300
atmosmodj	200	200	300	300	300	300
cage13	30	30	30	30	45	45
cage14	30	30	30	30	30	30
crankseg_1	6 300	6 200	6 300	6 300	7 300	5 900
CurlCurl_2	9 900	9 900	9 900	9 900	12 300	12 300
CurlCurl_4	21 100	21 100	21 100	21 100	26 400	26 400
ecology2	900	900	900	900	1 800	1 800
F1	29 200	29 200	29 200	29 200	-	-
FEM_3D_thermal2	300	300	300	300	-	-
G3_circuit	28 200	28 200	27 500	28 200	29 200	-
language	29	29	58	58	145	87
mc2depi	10 400	10 200	12 100	12 900	19 500	19 500
ns3Da	1 400	1 400	1 400	1 400	-	-
parabolic_fem	3 500	3 500	3 500	3 500	4 100	4 100
poisson3Db	300	300	300	300	-	-
rajat31	4 000	4 000	2 900	2 000	1 700	2 500
stomach	300	300	300	300	300	300
t2em	4 800	4 800	4 800	4 800	5 100	5 100
thermal2	21 100	28 500	25 600	28 500	-	-
tmt_unsym	500	500	500	500	500	500
torso2	80	80	80	80	-	-
torso3	200	200	200	200	300	300




Fig. 2. Relative performance of GMRES with a scalar Jacobi preconditioner and MGS (top) or CGSR (bottom).

TABLE 3
Inner Iteration Counts for GMRES Preconditioned With Jacobi

Matrix	Double		Preconditioner		Mixed		Single	
	MGS	CGSR	MGS	CGSR	MGS	CGSR	MGS	CGSR
af_0_k101	18 200	14 300	16 800	14 900	19 800	12 700	-	-
af_shell19	21 000	20 500	22 700	23 600	27 600	28 400	-	-
apache2	11 700	11 700	10 400	10 900	9 800	8 500	16 700	17 200
atmosmodj	200	200	300	300	300	300	300	300
bone010	14 600	19 000	19 100	21 500	25 600	28 200	-	-
Bump_2911	3 500	3 500	4 100	3 900	4 100	4 100	16 400	16 600
cage13	22	19	22	22	22	22	33	33
cage14	22	22	22	19	22	22	22	22
crankseg_1	800	800	800	800	800	800	1 300	1 200
CurlCurl_2	1 500	1 500	1 500	1 500	1 500	1 500	1 500	1 500
CurlCurl_4	1 900	1 900	1 900	1 900	1 900	1 900	1 900	1 900
ecology2	800	800	800	800	800	800	1 600	1 600
F1	3 600	3 600	3 300	3 600	3 600	3 800	-	-
FEM_3D_thermal2	60	60	60	60	60	60	-	-
G3_circuit	1 100	1 100	1 100	1 100	1 100	1 100	1 200	1 200
hood	3 900	3 900	4 100	4 100	4 100	4 000	-	-
language	29	29	58	58	58	58	145	87
mc2depi	11 000	10 600	13 200	-	12 600	12 200	-	-
ns3Da	14 300	14 400	14 800	14 600	15 000	14 900	-	-
parabolic_fem	3 600	3 600	3 700	3 700	3 700	3 700	4 400	4 400
poisson3Db	400	400	400	400	400	400	-	-
pwtk	13 500	18 400	15 700	16 600	17 700	20 200	-	-
rajat31	600	600	700	700	700	700	700	800
stomach	130	100	130	100	130	130	130	130
t2em	4 800	4 800	4 800	4 800	4 800	4 800	5 100	5 100
thermal2	21 400	25 300	25 400	25 500	22 700	25 500	29 800	-
tmt_unsym	500	500	500	500	500	500	500	500
torso2	56	47	56	56	56	56	-	-
torso3	134	100	100	100	134	134	134	134
venkat01	96	96	96	96	96	96	96	96

First, the results when no preconditioner is used are shown in Fig. 1. The average speedups for the mixed-precision approach were 18 percent for MGS and 61 percent for CGSR. Furthermore, it provided a speedup for most of the tested matrices and with CGSR almost doubled the performance on many of the matrices. It only reduced performance for *atmosmodj* using MGS and *language*. The single-precision implementation only satisfied the target accuracy on 16 of the 23 problems; for these problems, it had average speedups of 0 and 35 percent, respectively. The total number of inner iterations was mostly consistent between the double- and mixed-precision implementations, as is shown in Table 2, with the single-precision implementations needing similar or more iterations. The most notable exception is *rajat31*, which converged in significantly fewer iterations with the mixed- and single-precision implementations, contributing to its much higher speedups. We have been unable to determine the source of this behavior; we speculate that the floating-point error happens to perturb the Krylov subspace to better contain the solution. For *language*, the baseline implementation converged after 29 iterations without restarting, whereas the mixed-precision implementation restarted once after 29 iterations for a total of 58 iterations. Of the matrices with provided right-hand sides, *atmosmodj* had a significant increase in the number of iterations, but mixed-precision was still able to

achieve a speedup with CGSR; *thermal2* also had an increase in iterations for MGS, but still performed on par with the baseline. For the matrices with comparable iteration counts, many can achieve approximately a $2\times$ speedup with CGSR; however, some only obtained modest speedups. These reduced speedups correlate with the matrices with a high number of nonzeros per row relative to the size of the basis; thus, transferring the column-index array of the matrix will have a larger influence on the runtime. As per Table 1, even matrices with condition larger than the inverse of single-precision unit roundoff could be solved more efficiently with the mixed-precision approach; although, this may depend on the right-hand side. Finally, the matrices with right-hand sides from the SuiteSparse collection overall behaved similar to those with generated right-hand sides.

Second are the results for a scalar Jacobi preconditioner, shown in Fig. 2. The average speedups for the mixed-precision implementation were 12 and 50 percent for MGS and CGSR, respectively. The single-precision preconditioner outperformed the double-precision implementation in some cases. However, it failed on *mc2depi* for CGSR and provided an average slowdown of 8 percent for both orthogonalization schemes on the successful matrices. The single-precision implementation failed to satisfy the target accuracy in 11 of the 30 problems; the remaining problems had average speedups of -8% (i.e., a




Fig. 3. Relative performance of CGSR GMRES with an ILU(0) preconditioner and MGS (top) or CGSR (bottom).

slowdown) and 23 percent, respectively. For MGS, there were five matrices where the mixed-precision implementation failed to outperform double precision. For CGSR, mixed-precision outperformed double-precision for all matrices except *language*. The total number of inner iterations for a matrix was mostly consistent between the tested configurations, as is shown in Table 3, with *language* again taking twice the iterations and *pwtk* taking fewer iterations with MGS compared to CGSR. Like the non-preconditioned results, better speedups were achieved on matrices that had a small number of non-zeros per row relative to the size of the basis, which relates to the costs of matrix-vector products compared to orthogonalization.

Third are the results for an ILU(0) preconditioner, shown in Fig. 3. The average speedup for mixed-precision was -9% and -7% for MGS and CGSR, respectively (i.e., a slowdown). For the single-precision preconditioner, the speedups were -8% and -9% instead. There are a few factors that we attribute the lack of improvement to. The single-precision implementation failed to produce an accurate enough solution to 7 of the problems; the remaining problems had average speedups of -11% and -10% , respectively. First, both sparse triangular solves have limited parallelism for the GPU to exploit, particularly when there are few nonzeros per row; this results in the GPU bandwidth being underutilized and limited benefit by reducing the size of the data. Furthermore, the poor performance of the triangular solves causes them to make up a large part of the performance. Second, the

factorization is done in double-precision for both implementations, making it a fixed cost in the performance. Third, because of the effectiveness of the preconditioner, a high percentage of matrices can be solved without restarting in double-precision; however, the mixed-precision implementation must always restart at least once. These restarts incur overhead by computing the solution update and new residual, and reduce the rate of convergence for some matrices, increasing the iteration count, by discarding information on the old Krylov subspace and interfering with GMRES's superlinear convergence [35]. Table 4 shows the relevant iteration counts and that the baseline can converge without restarting for 11 out of the 29 matrices.

Finally, are the results for an ILU(0) preconditioner with five Jacobi iterations for triangular solves, shown in Fig. 4. The speedup was 8 and 13 percent for MGS and CGSR, respectively. Additionally, the single-precision preconditioner was able to achieve some improvement overall, with speedups of 3 and 4 percent, respectively. The single-precision implementation failed to produce an accurate enough solution to 3 of the problems; the remaining problems had average speedups of 4 and 4 percent, respectively. Note that while the triangular solves have been improved, the other factors limiting the improvement of the regular ILU(0) preconditioner remain. Table 5 shows the iteration counts and that the baseline only needed to restart on 5 of the 13 matrices.

While testing the performance, we discovered that using CGSR provides better performance in GPU-

TABLE 4
Inner Iteration Counts for GMRES Preconditioned With ILU(0)

Matrix	Double		Preconditioner		Mixed		Single	
	MGS	CGSR	MGS	CGSR	MGS	CGSR	MGS	CGSR
af_0_k101	4 500	4 400	5 500	7 700	5 200	6 600	7 700	7 800
af_shell19	2 200	2 200	2 200	2 200	2 400	2 500	-	-
apache2	600	600	600	600	800	800	700	700
atmosmodj	82	82	82	82	200	164	200	164
BenElechi1	3 900	4 300	4 600	4 500	4 200	4 500	-	-
bone010	1 200	1 200	1 200	1 200	1 200	1 200	2 200	4 700
cage13	7	7	8	8	8	8	12	12
cage14	7	7	7	7	8	8	8	8
crankseg_1	200	200	200	200	200	200	400	300
CurlCurl_2	600	600	600	600	600	600	700	700
CurlCurl_4	1 400	1 400	1 400	1 400	1 400	1 400	1 700	1 700
ecology2	200	200	200	200	200	200	300	300
F1	1 000	1 000	1 000	1 000	1 000	1 000	-	-
FEM_3D_thermal2	10	10	12	12	12	12	-	-
G3_circuit	200	200	200	200	300	300	300	300
language	9	9	14	14	14	14	28	42
marine1	300	300	300	300	300	300	300	300
mc2depi	1 700	1 700	1 600	1 600	1 700	1 700	1 800	1 700
ns3Da	200	200	200	200	200	200	-	-
parabolic_fem	800	800	800	800	800	800	900	900
poisson3Db	100	100	174	174	174	174	-	-
rajat31	10	10	9	9	18	18	18	18
stomach	17	17	18	18	20	18	20	18
t2em	600	600	600	600	600	600	600	600
thermal2	4 800	5 100	5 100	5 100	5 200	5 100	5 700	5 700
tmt_unsym	200	200	200	200	200	200	200	200
torso2	11	11	12	12	12	12	-	-
torso3	35	35	48	48	48	48	48	48
venkat01	12	12	12	12	12	12	12	12

accelerated GMRES compared to MGS, despite the extra orthogonalization. However, results showing this performance difference do not appear in literature, outside of a few assertions that CGSR is better for GPU-accelerated systems [19], [41]. The overall speedup is higher for simpler preconditioners and the mixed-precision implementation, as shown in Table 6. Recall that MGS requires j dot-products alternated with j vector additions for the j th inner iteration while CGSR merely requires four matrix-vector products. Thus, CGSR launches significantly fewer

kernels which reduces overhead; furthermore, high GPU utilization is easier to obtain with larger kernels than smaller ones. The better speedup for mixed- and single-precision implementations likely comes from reductions in the cost of the kernel’s execution making the kernel launches more costly relative to the total time. Similarly, when GMRES uses a cheaper preconditioner, it spends a higher percentage of its runtime doing the orthogonalization which results in a better speedup for switching from MGS to CGSR.




Fig. 4. Relative performance of CGSR GMRES with an ILU(0) preconditioner using five Jacobi iterations for triangular solves and MGS (top) or CGSR (bottom).

TABLE 5
Inner Iteration Counts for GMRES Preconditioned Using ILU(0) With Five Jacobi Iterations for Triangular Solves

Matrix	Double		Preconditioner		Mixed		Single	
	MGS	CGSR	MGS	CGSR	MGS	CGSR	MGS	CGSR
cage13	7	7	8	8	8	8	12	12
cage14	7	7	8	8	8	8	8	8
language	9	9	14	14	14	14	21	35
ns3Da	200	200	200	200	200	200	-	-
parabolic_fem	800	800	800	800	800	800	900	900
poisson3Db	100	100	174	174	174	174	-	-
stomach	21	21	22	22	22	22	22	22
t2em	800	800	800	800	800	800	800	800
thermal2	5 100	5 100	5 000	5 100	5 100	5 100	5 700	5 800
tmt_unsym	200	200	200	200	200	200	200	200
torso2	11	11	12	12	12	12	-	-
torso3	48	48	64	64	64	64	64	64
venkat01	16	16	16	16	16	16	16	16

TABLE 6
Average Speedup of CGSR-GMRES versus MGS-GMRES for Various Configurations

Preconditioner	Double	Mixed	Single
	Speedup	Speedup	Speedup
Identity	36%	87%	181%
Jacobi	34%	79%	116%
ILU(0)	4%	7%	4%
ILU(0) with Jacobi	8%	13%	4%

6 CONCLUSION

Like previous works with similar uses of precision [7], [18], [19], [20], our mixed-precision implementation never required more than twice the total inner iterations than the double-precision implementation, and usually much less than twice the total inner iterations when many restarts are needed. This reinforces the ideas provided by Section 3. Furthermore, looking at the matrices tested without a preconditioner, as listed in Table 1, none of the matrices satisfy $n^2u < 1$, let alone $\mathcal{O}(n^2j^3)_{UK}(\overline{AV_j}) < 1$. So, $c_4(n, j)$ from Theorem 1 can likely be significantly improved, which correlates with the analysis of the types of dot-product bounds used.

Between the theoretical results in Section 3 and the experimental results in Section 5, there is strong evidence that this mixed-precision approach for GMRES retains double-precision accuracy. Performance improvement was less clear cut, with the ILU(0) preconditioner seeing a slowdown. However, tests with GPU-friendly preconditioners and baselines that restarted consistently showed speedups, especially with CGSR orthogonalization.

There are three main future directions for this work. The first direction is to understand the performance of mixed-precision GMRES on multi-GPU and distributed systems. These systems are important for solving problems too large to be solved effectively or even fit on a single compute unit. However, they have additional communication costs to coordinate and exchange data. The second direction is to investigate the use of alternative data representations. 16-bit floating-point formats are one possibility but reduce the

accuracy by half compared to single precision. However, alternative techniques, such as compression, may allow using less than 32-bits per value without significantly reducing accuracy. Furthermore, it may be possible to reduce the memory needed for matrix indices. The third direction is to extend these ideas to other formulations of GMRES and other Krylov methods. One particularly important class of GMRES variants are the communication avoiding and pipelined algorithms, which try to reduce communication overheads when running on distributed systems.

ACKNOWLEDGMENTS

This work was supported in part by the University of Tennessee under Grant MSE E01-1315-038 as Interdisciplinary Seed funding, in part by UT Battelle subaward under Grant 4000123266 and in part by National Science Foundation under OAC under Grant 2004541.

REFERENCES

- [1] Y. Saad and M. H. Schultz, "GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems," *SIAM J. Sci. Stat. Comput.*, vol. 7, no. 3, pp. 856–869, 1986.
- [2] W. E. Arnoldi, "The principle of minimized iteration in the solution of the matrix eigenvalue problem," *Quart. Appl. Math.*, vol. 9, pp. 17–29, 1951.
- [3] Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd ed. Philadelphia, PA, USA: SIAM Press, 2003.
- [4] C. C. Paige and Z. Strakos, "Residual and backward error bounds in minimum residual Krylov subspace methods," *SIAM J. Sci. Comput.*, vol. 23, no. 6, pp. 1898–1923, May 2002.
- [5] L. Giraud, J. Langou, and M. Rozložník, "The loss of orthogonality in the Gram-Schmidt orthogonalization process," *Comput. Math. Appl.*, vol. 50, no. 7, pp. 1069–1075, Oct. 2005.
- [6] L. Giraud, J. Langou, M. Rozložník, and J. van den Eshof, "Rounding error analysis of the classical Gram-Schmidt orthogonalization process," *Numerische Mathematik*, vol. 101, no. 1, pp. 87–100, July 2005.
- [7] N. Lindquist, P. Luszczek, and J. Dongarra, "Improving the performance of the GMRES method using mixed-precision techniques," in *Proc. Driving Sci. Eng. Discov. Convergence HPC, Big Data AI*, 2020, pp. 51–66.
- [8] A. Abdelfattah et al., "A survey of numerical linear algebra methods utilizing mixed-precision arithmetic," *Int. J. High Perform. Comput. Appl.*, vol. 35, no. 4, pp. 344–369, Mar. 2021.
- [9] J. H. Wilkinson, *Rounding Errors in Algebraic Processes*. Princeton, NJ, USA: Prentice-Hall, 1963.

- [10] A. Buttari, J. Dongarra, J. Langou, J. Langou, P. Luszczyk, and J. Kurzak, "Mixed precision iterative refinement techniques for the solution of dense linear systems," *Int. J. High Perform. Comput. Appl.*, vol. 21, no. 4, pp. 457–466, Nov. 2007.
- [11] L. Giraud, A. Haidar, and L. T. Watson, "Mixed-precision preconditioners in parallel domain decomposition solvers," in *Proc. Domain Decomposition Methods Sci. Eng. XVII.*, 2008, pp. 357–364.
- [12] M. Baboulin et al., "Accelerating scientific computations with mixed precision algorithms," vol. 180, pp. 2526–2533, 2008.
- [13] S. Gratton, E. Simon, D. Tittley-Peloquin, and P. Toint, "Exploiting variable precision in GMRES," *SIAM J. Sci. Comput.*, 2020, *arXiv:1907.10550*.
- [14] J. V. D. Eshof and G. L. G. Sleijpen, "Inexact Krylov subspace methods for linear systems," *SIAM J. Matrix Anal. Appl.*, vol. 26, no. 1, pp. 125–153, Jan. 2005.
- [15] A. Bouras and V. Frayssé, "Inexact matrix-vector products in Krylov methods for solving linear systems: A relaxation strategy," *SIAM J. Matrix Anal. Appl.*, vol. 26, no. 3, pp. 660–678, Jan. 2005.
- [16] V. Simoncini and D. B. Szyld, "Theory of inexact Krylov subspace methods and applications to scientific computing," *SIAM J. Sci. Comput.*, vol. 25, no. 2, pp. 454–477, Jan. 2003.
- [17] H. Anzt, V. Heuveline, and B. Rucker, "Mixed precision iterative refinement methods for linear systems: Convergence analysis based on Krylov subspace methods," in *Proc. 10th Int. Conf. Appl. Parallel Sci. Comput.*, 2010, pp. 237–247.
- [18] H. Anzt, V. Heuveline, and B. Rucker, "An error correction solver for linear systems: Evaluation of mixed precision implementations," in *Proc. High Perform. Comput. Comput. Sci.*, 2011, pp. 58–70.
- [19] J. I. Aliaga, H. Anzt, T. Grützmacher, E. S. Quintana-Ort í, and A. E. Tomás, "Compressed basis GMRES on high performance GPUs," 2020, *arXiv:2009.12101*.
- [20] T. Iwashita, K. Suzuki, and T. Fukaya, "An integer arithmetic-based sparse linear solver using a GMRES method and iterative refinement," in *Proc. IEEE/ACM 11th Workshop Latest Adv. Scalable Algorithms Large-Scale Syst.*, 2020, pp. 1–8.
- [21] E. Agullo, F. Cappello, S. Di, L. Giraud, X. Liang, and N. Schenkels, "Exploring variable accuracy storage through lossy compression techniques in numerical linear algebra: A first application to flexible GMRES," Inria Bordeaux Sud-Ouest, Paris, France, *Res. Rep. RR-9342*, 2020.
- [22] A. Buttari, J. Dongarra, J. Kurzak, P. Luszczyk, and S. Tomov, "Using mixed precision for sparse matrix computations to enhance the performance while achieving 64-bit accuracy," *ACM Trans. Math. Softw.*, vol. 34, no. 4, pp. 17:1–17:22, Jul. 2008.
- [23] M. Emans and A. van der Meer, "Mixed-precision AMG as linear equation solver for definite systems," *Procedia Comput. Sci.*, vol. 1, no. 1, pp. 175–183, 2010.
- [24] M. Clark, R. Babich, K. Barros, R. Brower, and C. Rebba, "Solving lattice QCD systems of equations using mixed precision solvers on GPUs," *Comput. Phys. Commun.*, vol. 181, no. 9, pp. 1517–1528, 2010.
- [25] R. Strzodka and D. Goddeke, "Pipelined mixed precision algorithms on FPGAs for fast and accurate PDE solvers from low precision components," in *Proc. 14th Annu. IEEE Symp. Field-Programmable Custom Comput. Mach.*, 2006, pp. 259–270.
- [26] H. Anzt, G. Flegar, T. Grützmacher, and E. S. Quintana-Ort í, "Toward a modular precision ecosystem for high-performance computing," *Int. J. High Perform. Comput. Appl.*, vol. 33, no. 6, pp. 1069–1078, Nov. 2019.
- [27] O. S. Lawlor, "In-memory data compression for sparse matrices," in *Proc. 3rd Workshop Irregular Appl. Architectures Algorithms*, Nov. 2013, pp. 1–6.
- [28] N. Lindquist, "Reducing memory access latencies using data compression in sparse, iterative linear solvers," Bachelor's thesis, Saint John's Univ., New York, NY, USA, Apr. 2019.
- [29] C. C. Paige, M. Rozložník, and Z. Strakoš, "Modified Gram-Schmidt (MGS), least squares, and backward stability of MGS-GMRES," *SIAM J. Matrix Anal. Appl.*, vol. 28, no. 1, pp. 264–284, 2006.
- [30] J. Drkošová, A. Greenbaum, M. Rozložník, and Z. Strakoš, "Numerical stability of GMRES," *BIT Numer. Math.*, vol. 35, no. 3, pp. 309–330, Sept. 1995.
- [31] E. Carson and N. J. Higham, "Accelerating the solution of linear systems by iterative refinement in three precisions," *SIAM J. Sci. Comput.*, vol. 40, no. 2, pp. A817–A847, Jan. 2018.
- [32] N. J. Higham and T. Mary, "A new approach to probabilistic rounding error analysis," *SIAM J. Sci. Comput.*, vol. 41, no. 5, pp. A2815–A2835, Jan. 2019.
- [33] N. J. Higham, *Accuracy and Stability of Numerical Algorithms*, 2nd ed. Philadelphia, PA, USA: SIAM, 2002.
- [34] H. Anzt, J. Dongarra, and E. S. Quintana-Ortí, "Adaptive precision solvers for sparse linear systems," in *Proc. 3rd Int. Workshop Energy Efficient Supercomput.*, 2015, pp. 2:1–2:10.
- [35] H. A. Van der Vorst and C. Vuik, "The superlinear convergence behaviour of GMRES," *J. Comput. Appl. Math.*, vol. 48, no. 3, pp. 327–341, 1993.
- [36] H. C. Edwards, C. R. Trott, and D. Sunderland, "Kokkos: Enabling manycore performance portability through polymorphic memory access patterns," *J. Parallel Distrib. Comput.*, vol. 74, no. 12, pp. 3202–3216, 2014.
- [37] T. A. Davis and Y. Hu, "The University of Florida sparse matrix collection," *ACM Trans. Math. Softw.*, vol. 38, no. 1, pp. 1–25, Nov. 2011.
- [38] H. Avron, A. Druinsky, and S. Toledo, "Spectral condition-number estimation of large sparse matrices," *Numer. Linear Algebra Appl.*, vol. 26, no. 3, May 2019, Art. no. e2235.
- [39] H. Anzt, E. Chow, and J. Dongarra, "Iterative sparse triangular solves for preconditioning," in *Proc. Eur. Conf. Parallel Process.*, 2015, pp. 650–661.
- [40] E. Chow, H. Anzt, J. Scott, and J. Dongarra, "Using Jacobi iterations and blocking for solving sparse triangular systems in incomplete factorization preconditioning," *J. Parallel Distrib. Comput.*, vol. 119, pp. 219–230, Sep. 2018.
- [41] J. Dubois, C. Calvin, and S. Petiton, "Performance and numerical accuracy evaluation of heterogeneous multicore systems for Krylov orthogonal basis computation," in *Proc. Int. Conf. High Perform. Comput. Comput. Sci.*, 2011, pp. 45–57.



Neil Lindquist received the BA degree in computer science and mathematics from Saint John's University, Collegeville, Minnesota. He is currently working toward his graduate degree with Innovative Computing Laboratory, University of Tennessee, Knoxville's Tickle College of Engineering. His research interests include numerical linear algebra, and the effects of data representation on performance and accuracy.



Piotr Luszczyk received the BS and MSc degrees in computer science from the AGH University of Science and Technology in Kraków, Poland, and the PhD degree in computer science from the University of Tennessee Knoxville. He is currently a research assistant professor with Innovative Computing Laboratory, University of Tennessee, Knoxville's Tickle College of Engineering. His research interests include benchmarking, numerical linear algebra for high-performance computing, automatic performance tuning for modern hardware, and stochastic models for performance. He has more than a decade of experience developing HPC numerical software for large scale, distributed memory systems with multicore processors and hardware accelerators. He is currently a co-PI for the ECP xSDK project, the primary goal of which is to improve access to world-class software on exascale machines. He was the recipient of Google Scholar h-index of 38 and the i10-index of 105.



Jack Dongarra (Fellow, IEEE) is currently an American University distinguished professor of computer science with Electrical Engineering and Computer Science Department, University of Tennessee, a distinguished research staff member with Computer Science and Mathematics Division, Oak Ridge National Laboratory, and an adjunct professor with Computer Science Department, Rice University. He is the Turing Fellow with the Schools of Computer Science and Mathematics, University of Manchester, and a Faculty Fellow with the Institute for Advanced Study, Texas A&M University. He is the founding director of Innovative Computing Laboratory.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.