



**Subject Areas:**

Mathematics of computing,  
Mathematical software performance,  
Computing methodologies

**Keywords:**

FP16 arithmetic, half precision  
arithmetic, mixed precision solvers,  
LU factorization, iterative refinement,  
GMRES, GPU computing, linear  
algebra

**Author for correspondence:**

Azzam Haidar

e-mail: [azzamhaidar@nvidia.com](mailto:azzamhaidar@nvidia.com)

Mixed-Precision Iterative  
Refinement using Tensor  
Cores on GPUs to Accelerate  
Solution of Linear Systems

Azzam Haidar<sup>1</sup> Harun Bayraktar<sup>1</sup>

Stanimire Tomov<sup>2</sup> Jack Dongarra<sup>2,3,4</sup>

Nicholas J. Higham<sup>4</sup>

<sup>1</sup>NVIDIA, USA

<sup>2</sup>Department of Electrical Engineering and Computer  
Science, University of Tennessee, Knoxville,  
Tennessee, USA

<sup>3</sup>Computer Science and Mathematics Division, Oak  
Ridge National Laboratory, Oak Ridge, Tennessee,  
USA

<sup>4</sup>Department of Mathematics, University of  
Manchester, Manchester, M13 9PL, UK

Double-precision floating-point arithmetic (FP64) has been the de facto standard for engineering and scientific simulations for several decades. Problem complexity and the sheer volume of data coming from various instruments and sensors motivate researchers to mix and match various approaches to optimize compute resources, including different levels of floating-point precision. In recent years, machine learning has motivated hardware support for half-precision floating-point arithmetic. A primary challenge in high-performance computing is to leverage reduced-precision and mixed-precision hardware. We show how the FP16/FP32 Tensor Cores on NVIDIA GPUs can be exploited to accelerate the solution of linear systems of equations  $Ax = b$  without sacrificing numerical stability. On the NVIDIA Quadro GV100 (Volta) GPU, we achieve a  $4\times$ – $5\times$  performance increase and  $5\times$  better energy efficiency versus the standard FP64 implementation while maintaining an FP64 level of numerical stability.

# Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
<b>2</b>	<b>Related Work</b>	<b>4</b>
<b>3</b>	<b>Contributions</b>	<b>4</b>
<b>4</b>	<b>Iterative Refinement Solver: Background</b>	<b>5</b>
	(a) Iterative Refinement	5
	(b) The LU factorization	6
<b>5</b>	<b>Multi-Precision Factorizations</b>	<b>6</b>
	(a) Introducing Auto-Adaptive Rounding Technique	8
<b>6</b>	<b>Iterative Refinement Using Preconditioned GMRES for the Correction Equation IRGM</b>	<b>8</b>
<b>7</b>	<b>Preconditioned GMRES</b>	<b>9</b>
<b>8</b>	<b>Scaling Techniques</b>	<b>10</b>
	(a) Scalar Scaling	10
	(b) Two-Sided Diagonal Scaling	11
	(c) Two-Sided Diagonal Scaling Followed by Scalar Scaling	11
<b>9</b>	<b>Multiple Right-Hand Side Optimizations</b>	<b>11</b>
<b>10</b>	<b>Sensitivity of Performance to FP64 Compute Throughput</b>	<b>12</b>
<b>11</b>	<b>Energy Efficient Implementation</b>	<b>12</b>
<b>12</b>	<b>Performance Analysis</b>	<b>13</b>
<b>13</b>	<b>Experimental setup</b>	<b>14</b>
<b>14</b>	<b>Numerical Behavior</b>	<b>15</b>
	(a) Convergence Behavior	15
	(b) Scaling Techniques	19
<b>15</b>	<b>Performance</b>	<b>21</b>
	(a) Single Right-Hand Sides	22
	(b) Multiple Right-Hand Sides	25
	(c) Sensitivity to FP64 Compute Throughput	26
<b>16</b>	<b>Energy Efficiency</b>	<b>26</b>
<b>17</b>	<b>Conclusions and Future Directions</b>	<b>27</b>

## 1 Introduction

A fundamental requirement in scientific computing is the ability to solve a system of linear equations

$$Ax = b, \quad (1.1)$$

where  $A$  is a large, dense  $n \times n$  nonsingular matrix. This can be done at a speed that is close to the peak performance on current computer architectures, for example, by using libraries such as the NVIDIA CUSOLVER library [35], MAGMA [27,46], and MKL [28] that redesign and highly optimize the

standard LAPACK algorithms [1] for GPU and multi-core architectures. The solvers use direct methods in a fixed/working precision arithmetic, namely, the IEEE standard double-precision 64-bit floating-point arithmetic (FP64), or single precision 32-bit floating-point arithmetic (FP32). Recently, various machine learning and artificial intelligence neural network applications increased the need for half precision arithmetic, and vendors started to accelerate it in hardware, in the form of either the IEEE FP16 format [26] or the bfloat16 format [29] (see Table 1.1). Currently, the NVIDIA V100 Tensor Cores (TCs) can execute FP16-TC at up to 120 teraFLOP/s versus 7.5 teraFLOP/s for FP64 and 15 teraFLOP/s for FP32. Thus, developing algorithms to exploit the much higher performance that lower-precision arithmetic offers can have significant impact in scientific, high-performance computing (HPC).

This paper presents a class of mixed-precision algorithms and an accompanying set of computational techniques that we have developed to accelerate (1.1) in FP64, which is the de facto standard for scientific computing. We show that the new mixed-precision techniques can accelerate the solution by a significant factor using the faster lower precisions, while still retaining FP64 quality.

The mixed-precision approach is based on an iterative refinement process in FP64 arithmetic:

$$x_{i+1} = x_i + c_i, \quad (1.2)$$

which is performed until  $x_{i+1}$  has a backward error at the FP64 level, or some other user-specified tolerance. Here,  $c_i$  is a correction vector obtained by mixed-precision solution of the correction equation  $Ac_i = b - Ax_i$ .

The fundamentals of the mixed-precision approach (1.2) are based on the mixed-precision iterative refinement algorithm. The main idea is to compute an LU factorization of  $A$  in low precision and use the factorization to solve the correction equation. These methods have been studied in the past, as discussed in Section 2. A persistent challenge has been to redesign the techniques for new architectures and to develop highly tuned implementations that resolve computational issues such as inefficient parallelization, scaling, and use of mixed-precision calculations. To address this problem on GPU Tensor Cores, we develop a number of innovations for mixed-precision computations (outlined in Section 3) as well as leverage building blocks from HPC numerical libraries such as CUSOLVER and MAGMA, which provide state-of-the-art, high-performance algorithms such as LU factorization—including a set of highly tuned mixed-precision iterative refinement algorithms using either the FP32 or the FP16 as lower precision for the LU factorization (e.g., FP32  $\rightarrow$  FP64 and FP16  $\rightarrow$  FP64) [17,45].

Recently, some work on the HPL-AI [12] benchmark—which is a benchmark similar to the well-known HPL benchmark [25] that is used for ranking the TOP500 supercomputers in the world, but which uses iterative refinement to solve the linear system—showed unprecedented and attractive results on the top machine in the world. The HPL-AI benchmark using iterative refinement with FP16-TC as the lower precision achieved 445 petaFLOP/s (June 2019) and 550 petaFLOP/s (November 2019) on the Summit supercomputer [37] while the FP64 standard HPL scored 148.6 petaFLOP/s [5,38].

Table 1.1: Parameters for the IEEE FP16, FP32, and FP64 arithmetic precisions, the mixed-precision FP16-TC, and their respective peak performances on an NVIDIA V100 GPU. “Range” denotes the order of magnitude of the smallest subnormal ( $x_{\min,s}$ ), and largest and smallest positive normalized floating-point numbers. The FP16-TC features mixed-precision arithmetic operations, where inputs are FP16 and computations are performed in full (FP32) precision.

Arithmetic	Size (bits)	Range			Unit roundoff	Peak Tflop/s (V100 GPU)
		$x_{\min,s}$	$x_{\min}$	$x_{\max}$		
FP16-TC	16/32					120
BFloat16	16		$1.2 \times 10^{-38}$	$3.4 \times 10^{38}$	$3.9 \times 10^{-3}$	N/A
FP16	16	$6.0 \times 10^{-8}$	$6.1 \times 10^{-5}$	$6.6 \times 10^4$	$4.9 \times 10^{-4}$	30
FP32	32	$1.4 \times 10^{-45}$	$1.2 \times 10^{-38}$	$3.4 \times 10^{38}$	$6.0 \times 10^{-8}$	15
FP64	64	$4.9 \times 10^{-324}$	$2.2 \times 10^{-308}$	$1.8 \times 10^{308}$	$1.1 \times 10^{-16}$	7.5

## 2 Related Work

Iterative refinement is a long-standing method that was programmed by Wilkinson in the 1940s for the ACE digital computer. The idea is to improve the computed solution of a linear system by iteratively solving a correction equation and adding the correction to the current solution; see Wilkinson [47], Moler [34], Stewart [43], Demmel [11], and, for a comprehensive treatment, Higham [21, Chap. 12]. The three tasks—original solve/factorization, residual computation, and correction equation solve—can be done in the same precision (fixed precision) or in different precisions (mixed precision). The original usage was mixed precision, with the residual computed at twice the working precision. Fixed-precision iterative refinement was analyzed by Skeel [42] for a solver based on LU factorization and by Higham [19,20] for a general solver. In the 2000s, motivated by processors equipped with FP32 that had speed  $2\times$  that of FP64, mixed-precision iterative refinement—with the LU factorization done in FP32 and everything else done in FP64—was explored in [6,31].

Replacing the direct triangular solves of the correction equation with an iterative method, as suggested in [7] in a mixed-precision context, leads to “nesting” of two iterative methods, which in general are called “inner–outer” iterations, the latter having been studied both theoretically and computationally [15,39,41], including in mixed-precision computation scenarios [3]. Recently, Carson and Higham [7,8] analyzed a three-precision iterative refinement scheme (factorization precision, working precision, residual precision) and concluded that if the condition number of  $A$  is not too large, namely  $\kappa_\infty(A) = \|A\|_\infty \|A^{-1}\|_\infty < 10^4$ , then using FP16 for the  $O(n^3)$  portion (the LU factorization) and (FP32, FP64) or (FP64, FP128) as the (working, residual) precision for the  $O(n^2)$  portion (refinement loop), one can expect to achieve forward error and backward error on the order of  $10^{-8}$  and  $10^{-16}$  respectively. We note that, if  $\hat{x}$  is an approximate solution of  $Ax = b$  the forward error is defined by  $\|\hat{x} - x\|_\infty / \|x\|_\infty$  and the backward error is defined by  $\min\{\varepsilon : (A + \Delta A)x = b, \|\Delta A\| \leq \varepsilon \|A\|\}$  and can be evaluated as  $\|r\|_2 / (\|A\|_2 \|\hat{x}\|_2)$ , where  $r = b - A\hat{x}$ . Carson and Higham also proposed the use of the GMRES method preconditioned by the FP16 LU factorization as the refinement procedure and showed that in this case the constraint on the condition number can be relaxed to  $\kappa_\infty(A) < 10^8$  when the (working, residual) precision is (FP32, FP64) and to  $10^{12}$  when the (working, residual) precision is (FP64, FP128). Analysis covering this GMRES-based approach when two precisions are used with the residual precision equal to the working precision is given in [22].

An investigation of similar iterative refinement methods on earlier generations of GPUs can be found in [18].

## 3 Contributions

The primary motivation of this paper is to propose and implement a high-performance framework for the mixed-precision iterative refinement solvers that make use of GPU Tensor Core–accelerated FP16–TC. To this end, we:

- Develop a framework for exploiting GPU Tensor Cores in mixed-precision (FP16–TC–FP32/FP64) iterative refinement solvers and describe the path to develop high-performance, Tensor-Cores-enabled dense linear algebra building blocks that can be used to exploit the FP16–TC in HPC applications.
- Provide algorithmic advancement to minimize trouble related to matrices not representable in FP16’s limited range. Matrices from a wide range of applications become singular in FP16 because their entries are not representable in the limited FP16 range. In our algorithm, the original data is never represented in precision below FP32. This technique is the first step toward our proposition for a new set of mixed-precision factorization algorithms. In addition, we investigate scaling techniques that can also help fix issues related to the FP16 range.
- Introduce a new class of multi-precision dense matrix factorization algorithms. Historically, iterative refinement solvers have always used only one precision for the factorization (lower precision) and a higher (one or two) precision for the refinement loop. Such a design means that the factorization will be performed in FP16, the refinement loop in FP64. However, as our study will show, performing the whole factorization in FP16 is not optimal, for two main reasons [24].

First, an FP16 factorization has reduced accuracy. Second, having to represent the original matrix in FP16 can be problematic. In our work, the factorization itself is implemented in multiprecision to allow better accuracy and to be able to solve a wider range of problems than was possible before.

- Compare iterative refinement with preconditioned GMRES as the solver with preconditioned GMRES applied directly to the linear system (with no iterative refinement).
- Develop a performance model that accurately predicts performance gains, allowing users to decide in advance if such iterative refinement solvers can be beneficial for their problems.
- Present experiments that show that many problems, representing different areas of applications, with both dense and sparse matrices, can be accelerated by our mixed-precision iterative refinement solver up to  $5\times$  when using the FP16-TC, or about  $2\times$  when using the FP32 arithmetic.
- Quantify—in practice—the performance and limitations of this approach on the latest hardware architecture from NVIDIA, namely, the Volta V100 GPU.
- Show how the performance of the mixed-precision iterative refinement solver is not sensitive to the FP64 compute throughput.
- Study the benefits and applicability of TCAIRS for linear systems with multiple right-hand sides.
- Provide experiments for both real and complex data.
- Present a study of the energy efficiency of the iterative refinement solver exploiting the Tensor Cores and show that it can reduce power consumption by up to 80%.

The developments are released within the CUDA toolkit CUSOLVER library [35] and in the MAGMA library [17,33] to make these experiments independently reproducible and to allow the scientific community to exploit this work.

We also would like to point readers who are interested in energy efficiency and power measurement to our work in [16] which investigates the energy gain that can be brought by the iterative refinement techniques from a power consumption point of view.

## 4 Iterative Refinement Solver: Background

The standard method for solving a linear system  $Ax = b$  is Gaussian elimination with partial pivoting, or equivalently, the LU factorization method with partial pivoting. An LU factorization represents  $A$  as the product of a lower triangular matrix  $L$  and an upper triangular matrix  $U$ , so that  $A = LU$ . Thus, solving  $Ax = b$  reduces to solving two triangular systems:

$$Ax = b \Rightarrow LUx = b : \text{ solve } Ly = b \text{ then solve } Ux = y.$$

The factorization costs  $O(n^3)$  operations while the triangular solve costs  $O(n^2)$  operations, so for large  $n$  almost all the time is spent in the factorization and the triangular solve time is negligible. In practice, partial pivoting is used to ensure numerical stability, resulting in a factorization  $PA = LU$ , where  $P$  is a permutation matrix. For simplicity of exposition, we suppress  $P$  in what follows (thus  $A \leftarrow PA$ ). Throughout the paper  $A$  is assumed to be nonsingular.

The basic way to solve such a system is to perform these operations in one precision, typically the precision to which the user wishes to obtain the solution, or the precision in which the input data is given. We call this the working precision and denote it by  $\mathbf{u}^w$ .

### (a) Iterative Refinement

Iterative refinement aims to improve the accuracy or backward error of a computed solution  $\hat{x}$  to  $Ax = b$ . It consists of a series of iterations (the refinement loop) and is described in Algorithm 4.1 in the general three-precision form proposed by Carson and Higham [8]. The three precisions of arithmetic used are defined by their respective unit roundoffs:  $\mathbf{u}^f \geq \mathbf{u}^f \geq \mathbf{u}^r$ . We write “in precision  $\mathbf{u}^f$ ” (etc.) to mean “in floating-point arithmetic of precision  $\mathbf{u}^f$ ”.

As seen in Algorithm 4.1, we denote by  $\mathbf{u}^r$  the precision in which step 1 is carried out. Step 2 is performed in the precision  $\mathbf{u}^f$ , which is the precision of the LU factorization: it uses the LU factors to solve the correction equation  $Ac = r$  and then it casts the solution  $c_i$  to the working precision  $\mathbf{u}^w$ . Step 3 is carried out with precision  $\mathbf{u}^w$ . If all three steps can be computed exactly then the iterative refinement

```

Convert  $A \rightarrow A^f$  from precision  $\mathbf{u}^w$  to  $\mathbf{u}^f$ ;
Perform LU factorization of  $A^f$  in precision  $\mathbf{u}^f$  ;
Find the initial solution  $x_0$  using the computed LU factorization of  $A^f$  in precision  $\mathbf{u}^f$ , then cast  $x_0$  to
precision  $\mathbf{u}^w$  ;
//Refinement loop, outer loop
repeat
  1. Residual: Compute residual  $r_i \leftarrow b - Ax_i$  in precision  $\mathbf{u}^f$  and cast it to  $\mathbf{u}^w$  ;
  2. Correction: Solve  $Ac_i = r_i$ , using the LU factorization in precision  $\mathbf{u}^f$  then cast the result  $c_i$  to the
working precision  $\mathbf{u}^w$  ;
  3. Update: Correct the current solution  $x_{i+1} = x_i + c_i$  in precision  $\mathbf{u}^w$ ;
until  $x_{i+1}$  is accurate enough;

```

**Algorithm 4.1:** Classical iterative refinement (IR) algorithm.

algorithm completes in one iteration. However, in floating-point arithmetic the above steps usually need to be repeated.

If  $\mathbf{u}^f = \mathbf{u}^w = \mathbf{u}^r$  the method is called fixed-precision iterative refinement; otherwise, it is mixed-precision iterative refinement. Fixed-precision iterative refinement can be used to improve the backward error of an LU factorization without a strong stabilizing pivoting strategy [2,32,42]. On the other hand, mixed-precision iterative refinement also improves the forward error to the working precision—if the condition number of  $A$  is not too large:  $\mathbf{u}^f \kappa_\infty(A) \leq 1$ . We denote this method by IR.

In our analysis and experiments we take  $\mathbf{u}^r = \mathbf{u}^w$ , so we do not use extra precision in forming the residuals.

## (b) The LU factorization

Algorithmically, as presented in Algorithm 4.2 and illustrated in Figure 4.1, the LU factorization can be viewed as a sequence of steps with two distinct phases per step: (1) a panel factorization that affects the data depicted by the orange portion of Figure 4.1, and 2) a trailing matrix update that updates data represented by the magenta (denoted by  $T_i$ ) and green colors ( $A_i$ ) in Figure 4.1. From a software point of view, we know that `PanelFactorize` is a memory-bound step performed through the `Xgetf2` routine and occupies a small portion of the total time, while `TrailingMatrixUpdates` are compute-bound and are performed using the Level-3 Basic Linear Algebra Subprograms (BLAS) routines `Xtrsm` (for  $T_i$ ) and `Xgemm` (for  $A_i$ ). The  $A_i$  updates occupy the greatest portion of the time spent in the factorization. Thus one might expect the performance of the LU factorization to be asymptotically similar to the Level-3 BLAS `Xgemm` routine.

```

for  $P_i \in \{P_1, P_2, \dots, P_n\}$  do
  PanelFactorize( $P_i$ )
  TrailingMatrixUpdate( $T^{(i)}$ )
  TrailingMatrixUpdate( $A^{(i)}$ )

```

**Algorithm 4.2:** The LU factorization process.

## 5 Multi-Precision Factorizations

In contrast to the standard factorization that works in one precision (see Algorithm 4.2), we have developed a new class of multi-precision factorizations that use a higher precision calculation  $\mathbf{u}^h$  for critical parts of the algorithm while using a lower precision  $\mathbf{u}^l$  for the other parts (the parts that are numerically less sensitive and also the most time consuming). The most useful case of the multi-precision factorization approach is when FP16 is used for the factorization, i.e., when  $\mathbf{u}^f$  of Algorithm 4.1 is equal to FP16.




Figure 4.1: Two-phase implementation of a one-sided factorization.

Originally, in our previous work [18], we found that performing the factorization in only FP16 suffers from many issues, such as the entries of the matrix  $A$  being outside the representable range of the FP16, making the matrix singular in  $\mathbf{u}^f$ , as well as causing numerical troubles during the panel computation. By using the multi-precision approach, we can maintain the benefit of the high speed of the FP16 precision while producing a better solution than performing the whole factorization in FP16.

For example, in the case of a FP32-FP16 LU factorization, the  $O(n^2)$  work in the panel factorization will be computed in FP32, and the  $O(n^3)$  FLOPs in the trailing matrix updates will be done in FP16. This is illustrated in Algorithm 5.1. Note that the trailing matrix update for  $A_i$  is a mixed-precision Xgemm, where  $P_i^{f_\ell}$  and  $T_i^{f_\ell}$  are of lower precision  $\mathbf{u}^{f_\ell}$ , while the input  $A_i^{f_h}$  and the resulting  $A_i^{f_h}$  are in higher (FP32) precision  $\mathbf{u}^{f_h}$ . This type of operation is a feature of the NVIDIA tensor cores hardware, and we are trying to take advantage of it by providing a fast factorization while maintaining the sensitive portion of the computation in higher precision. For the sake of completeness, we briefly describe how the NVIDIA Tensor Cores operate and allow us to get the  $A_i^{f_h}$  in higher (FP32) precision  $\mathbf{u}^{f_h}$ . The Tensor Cores are programmable matrix-multiply-and-accumulate units that can deliver up to 120 teraFLOPs/s on NVIDIA Volta GPU hardware. This number is higher on newer NVIDIA hardware. On an NVIDIA Volta GPU, each Tensor Core performs the operation  $D = A \times B + C$ , which is an ‘‘FP16 input multiply with full-precision product and FP32 accumulate’’ as shown in Figure 5.1. The matrix inputs  $A$  and  $B$  are in FP16, while the matrices  $C$  and  $D$  could be in FP16 or FP32. In our case the matrix  $D$  is the matrix  $C$  and it is in FP32.

We emphasize that the different precisions used here in the multi-precision factorization are not to be confused with the different precisions used in the refinement process. The  $\mathbf{u}^f$  precision denoted in Algorithm 4.1 can be viewed as an umbrella for  $\mathbf{u}^{f_h}$  and  $\mathbf{u}^{f_\ell}$ . The input and output data for the factorization is in  $\mathbf{u}^{f_h}$  and the factorization is carried out using both the  $\mathbf{u}^{f_h}$  and  $\mathbf{u}^{f_\ell}$  precisions. Then it becomes a notation preference to denote  $\mathbf{u}^f$  to be either of  $\mathbf{u}^{f_h}$  or  $\mathbf{u}^{f_\ell}$ —but no matter the notation, the factorization will be carried out with both precisions. This type of mixed-precision factorization was first developed and released within the MAGMA library [17,33] and the CUDA toolkit CUSOLVER library [35].

```

Input:  $A$  in precision  $\mathbf{u}^{f_h}$ 
 $\mathbf{u}^{f_\ell} \leq \mathbf{u}^{f_h}$ 
for  $P_i \in \{P_1, P_2, \dots, P_n\}$  do
    PanelFactorize( $P_i$ ) in precision  $\mathbf{u}^{f_h}$ 
    TrailingMatrixUpdate( $T_i$ ) in precision  $\mathbf{u}^{f_h}$ 
    convert  $P_i \rightarrow P_i^{f_\ell}$  from precision  $\mathbf{u}^{f_h}$  to  $\mathbf{u}^{f_\ell}$ 
    convert  $T_i \rightarrow T_i^{f_\ell}$  from precision  $\mathbf{u}^{f_h}$  to  $\mathbf{u}^{f_\ell}$ 
    TrailingMatrixUpdate( $A_i^{f_h} = A_i^{f_h} - P_i^{f_\ell} T_i^{f_\ell}$ ) using tensor cores

```

Algorithm 5.1: The multi-precision LU factorization process.

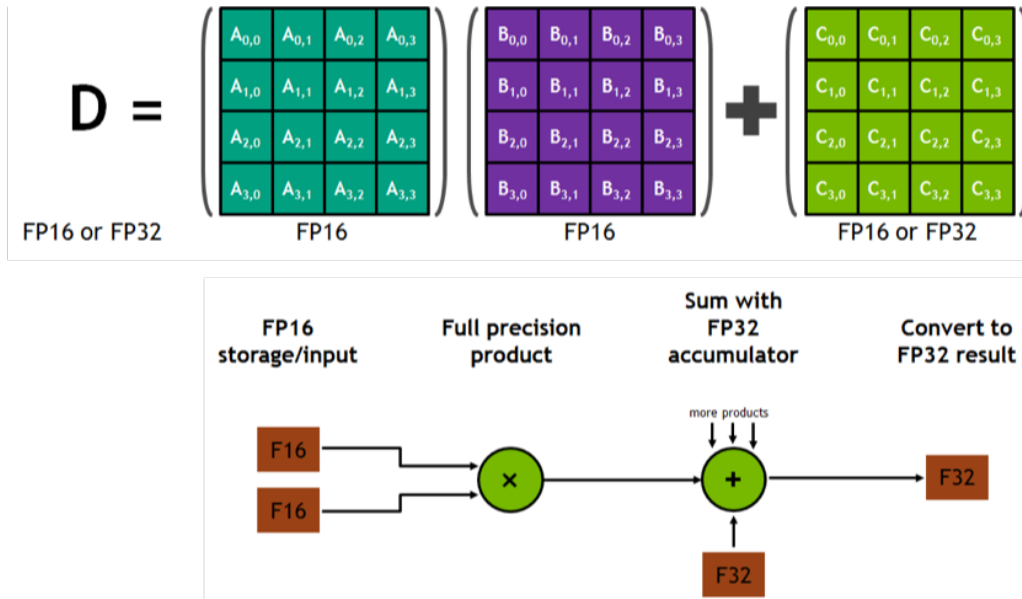


Figure 5.1: Tensor Core matrix multiply and accumulate, showing how products of scalars within a  $4 \times 4$  matrix multiplication are formed and accumulated.

The benefit of multi-precision factorization is two-fold. First, it helps solve many of the FP16 numerical issues, as the input matrix  $A$  is never converted to FP16, thus there are no issues related to the conversion of  $A$ . Second, all the sensitive portions of the computation (panel factorization and triangular solve) happen in FP32, thus there are no issues related to FP16 underflow or overflow in these portions. It is only during the trailing matrix update (e.g., Schur rank- $k$  update) of  $A^{(i)}$  where FP16 is involved, i.e., the use of a mixed-precision Xgemm of the form:

$$C_{\text{FP32}} = C_{\text{FP32}} - A_{\text{FP16}} B_{\text{FP16}}. \quad (5.1)$$

This means that issues related to the FP16 could only appear here where values of  $A_{\text{FP16}}$  and  $B_{\text{FP16}}$ , are large above the FP16 range. Thus we propose the auto-adaptive rounding mode to overcome this issue.

### (a) Introducing Auto-Adaptive Rounding Technique

We propose an adaptive rounding technique that rounds  $A$  and  $B$  by round to nearest, except for values larger than the FP16 range which are rounded to the nearer of the maximal or minimal normalized FP16 values (namely  $\pm 65504$ ). Thus our proposed adaptive rounding method avoids breakdown in the algorithm caused by overflows creating infinities.

As a result, our proposed multi-precision factorization ensures that during the factorization process as soon as the values are within the FP32 range, there will be no overflow related to the FP16. In practice, mixed-precision factorization gives accuracy between that of a fixed FP32 precision factorization and a fixed FP16 factorization. When the fixed FP16 precision factorization works, mixed precision can still produce a backward error between one and two orders of magnitude smaller than the pure FP16 factorization [4], while retaining the high performance of the FP16 variant.

## 6 Iterative Refinement Using Preconditioned GMRES for the Correction Equation IRGM

GMRES [40] is a popular Krylov subspace iteration for solving a general linear system of equations. Following Carson and Higham [7], we consider another variant of iterative refinement by using



preconditioned GMRES to approximately solve the correction equation  $Ac_i = r_i$  in step 2 of the classical IR in Algorithm 4.1. GMRES will be preconditioned by the low-precision LU factors. The idea is that the GMRES solver will provide a better and more stable approximate solution to  $Ac_i = r_i$  than the basic triangular solve, which is directly affected by the quality of the low-precision LU factors. Using GMRES, we can still guarantee that the solution of the correction equation  $Ac_i = r_i$  has some correct digits and a residual at the level of the convergence tolerance requested by the algorithm. The convergence tolerance is chosen to be of the order of the unit roundoff of the low-precision arithmetic used during the factorization (e.g., we use  $10^{-4}$  or  $10^{-8}$  when the LU factorization is in FP16 or FP32, respectively). We denote this variant by IRGM, and it is described in Algorithm 6.1. Note that  $U^{-1}$  and  $L^{-1}$  are never explicitly formed; instead matrix–vector products  $U^{-1}L^{-1}Ay$  needed by GMRES are computed by multiplication by  $A$  followed by two triangular solves. Since this paper focuses on the practical usage and possible performance gains rather than error analysis, we point the reader to [7,8,22] for detailed error analysis of the IR and IRGM techniques.

```

Convert  $A \rightarrow A^f$  from precision  $\mathbf{u}^w$  to  $\mathbf{u}^f$ ;
Perform LU factorization of  $A^f$  in precision  $\mathbf{u}^f$  ;
Find the initial solution  $x_0$  using the computed LU factorization of  $A^f$  in precision  $\mathbf{u}^f$  then cast  $x_0$  to
precision  $\mathbf{u}^w$  ;
//Refinement loop, outer loop
repeat
  1. Residual: Compute residual  $r_i \leftarrow b - Ax_i$  in precision  $\mathbf{u}^f$  and cast it to  $\mathbf{u}^w$ . ;
  2. Correction in IRGM : Solve  $U^{-1}L^{-1}Ac_i = U^{-1}L^{-1}r_i$  by GMRES in precision  $\mathbf{u}^w$  ;
  3. Update: Correct the current solution  $x_{i+1} = x_i + c_i$  in precision  $\mathbf{u}^w$ ;
until  $x_{(i)}$  is accurate enough;

```

**Algorithm 6.1:** iterative refinement using GMRES (IRGM) algorithm.

## 7 Preconditioned GMRES

The idea here is to use a preconditioned, full GMRES to solve the original linear system  $Ax = b$ , instead of using preconditioned GMRES to solve the correction equation as in Section 6. This idea can be viewed as setting the tolerance of the GMRES in Section 6 to the unit roundoff of FP64, thus the correction will be solved down to FP64 accuracy, meaning the refinement loop of Algorithm 6.1 will finish in one outer iteration. This proposition comes from the analysis of the Krylov subspace on which the GMRES method is based. In Section 6, the refinement loop can be viewed as a restarted GMRES, but instead of performing  $m$  iterations before each restart, it performs any number of iterations until it reaches the tolerance set for the GMRES (e.g.,  $10^{-4}$  in this study), while the method proposed here uses a full GMRES until convergence. It is well known that a full GMRES most likely converges faster than the restarted one because it uses all the previously constructed Krylov subspace to find the new direction to minimize the residual, while the restarted one starts constructing a new Krylov subspace after each restart. Our experiments show that for hard cases, where the convergence requires more than 30 iterations, the proposed preconditioned full GMRES has an advantage over the method proposed in Section 6, while for problems where the number of iterations is small, both methods behave the same within one to two iterations difference. In order to illustrate this behavior, we generate a hard test case and depict the convergence history in Figure 14.3 for both real and complex matrices. Since it is one of the hard cases, we can see that the classical IR cannot converge. The IRGM and GM refinement methods both follow the same path for the first iteration of the outer loop of the IRGM method, (e.g., the first 40 and 60 iterations for real and complex cases, respectively). After that the GM refinement method continues to decrease the residual using the entire previously generated Krylov subspace to find the new direction, while the IRGM method starts the construction of a new set of Krylov subspace vectors, resulting in slower convergence. We mention that for all other cases depicted in

Section 14 both the GM and the IRGM refinement methods behave similarly, while the IR method can fail sometimes. However, we caution that one does not normally run Krylov methods to convergence to the working precision and the convergence rate of GMRES is not well understood for general matrices  $A$ , so the behavior we have observed may not always hold.

## 8 Scaling Techniques

Numerical algorithms that deal with FP16 computations must deal with the limited range of representable elements in FP16 (see Table 1.1 for more details). For example, when converting a matrix  $A$  to  $A_{\text{FP16}}$ , the matrix entries may overflow or underflow. Our approach avoids overflow (as long as the FP32 range is not exceeded) and reduces the chance of underflow. Thanks to the proposed innovative multi-precision factorization  $A$  is never converted to FP16. Instead, it is converted to FP32 which, together with the use of auto-adaptive rounding method (see Section 5(a)) in the factorization, excludes such problems. Even though the proposed multi-precision factorization alleviates many of the FP16 numerical issues (e.g., related to conversion or to computation), some other issues still need attention. In particular, if the computation involves values close to or above the range limit then the quality of the LU factors might be compromised and might require other techniques to cope with the low bits that FP16 explores. One such technique is scaling. We note that, the bfloat16 arithmetic provide a range similar to the FP32 but with short mantissa and thus it could also be used for the factorization and avoids (as long as the FP32 range is not exceeded) overflow related to the FP16 range. However, we expect the short mantissa to be a limitation.

Originally, scaling was introduced in the context of iterative methods to speed up their convergence. Later, scaling was also used in direct methods (both dense and sparse) as a way to reduce the condition number of a matrix. In our work, we introduce scaling in the context of the iterative refinement solver. The goal of scaling is first to make a double- or single-precision matrix representable in FP16, as well as reduce its condition number since the convergence of the iterative refinement solver might be affected by the condition number of the matrix. Different approaches for scaling could be used. Research about scaling techniques for dense and sparse system can be found in [13,14]. More recently, [24] studied different scaling techniques for the iterative refinement solver using FP16 precision with tests on relatively small matrices. In this paper we study three approaches for scaling, similar to [24], and we show the effect of each method on the convergence of the iterative refinement solver. We first discuss a scalar scaling technique and then we discuss a two-sided diagonal scaling method. We also study a technique that uses a combination of the two scaling approaches.

### (a) Scalar Scaling

The idea here is to scale the matrix in such a way that its elements are mapped to be within a constant of the FP16 interval range, i.e.,  $a_{ij}^{(scaled)} \in [-\theta x_{\max}, \theta x_{\max}]$ , where  $\theta \in (0, 1]$  is a parameter. (Recall that  $x_{\max}$  is given in Table 1.1.) This technique is described in Algorithm 8.1. The strategy does two things: it ensures that the matrix fits into the FP16 range and it stretches the elements to cover a fraction  $\theta$  of the range, so as to make the best possible use of the limited range. Since the matrix is subsequently LU factorized, it is important that some headroom is allowed for growth during the factorization: if the elements grow by a factor  $\rho_n$  then  $\theta$  should be less than  $1/\rho_n$ . For many problems,  $\rho_n < 100$  with partial pivoting, so  $\theta = 0.01$  will be sufficient.

One of the effects of this algorithm is that it can result in many elements underflowing, which is possible if many of the elements of  $A$  are sufficiently smaller in magnitude than  $a_{\max}$ .

The effect of this scaling on the quality and the convergence of the iterative refinement solver is analyzed in Section 14.

1. Find  $a_{\max} = \max_{i,j} |a_{ij}|$
2. Let  $\mu = \theta x_{\max} / a_{\max}$
3. Scale  $A$  such that  $A^{(scaled)} = \mu A$

**Algorithm 8.1:** Scalar scaling.

## (b) Two-Sided Diagonal Scaling

To address the issues of the scalar scaling, let us study another scaling strategy that applies a two-sided diagonal scaling to  $A$  such that  $A^{(scaled)} = RAC$ , where  $R = \text{diag}(r_i)$  and  $C = \text{diag}(c_i)$  are diagonal matrices. Note that  $r_i$  and  $c_i, i = 1 : n$  are the scaling factors for each row and column of  $A$ , respectively. Such scaling algorithms have been developed in the context of linear systems, and in particular we focus on the scaling algorithm provided by LAPACK through the routine XGEEQU. We describe the two-sided diagonal scaling algorithm of XGEEQU in Algorithm 8.2. This technique computes row and column scaling intended to equilibrate a matrix  $A$  and reduce its condition number.  $R$  holds the row scale factors and  $C$  the column scale factors, chosen to make the largest element in each row and column of the matrix  $A^{(scaled)}$  (with elements  $A^{(scaled)}(i, j) = R(i) * A(i, j) * C(j)$ ) have an absolute value of 1 (note that every row and column must be nonzero since  $A$  is nonsingular).

```

for  $i \in \{1, 2, \dots, n\}$  do
   $R_i = 1 / \|A(i, :)\|_\infty$ 
 $\hat{A} = RA$ 
  //  $\hat{A}$  is row equilibrated
for  $j \in \{1, 2, \dots, n\}$  do
   $C_j = 1 / \|\hat{A}(:, j)\|_\infty$ 

```

**Algorithm 8.2:** Two-sided scaling.

## (c) Two-Sided Diagonal Scaling Followed by Scalar Scaling

The idea here is to use a combination of the scaling techniques described above. We represent the implementation in Algorithm 8.3. First, a two-sided scaling will be performed on  $A$  to smooth its values and minimize its condition number by making all the values within the FP16 range. Then, a scalar scaling will be applied (see lines 2–4 from Algorithm 8.3) in order to move the elements towards the largest representable number so as to make the best possible use of the limited FP16 range. More details about this technique can be found in [24]. We note that since the intent of the two-sided diagonal scaling is to make the maximal absolute value of  $\tilde{A}$  less than or equal to one, we can say that  $\beta$  in line 2 of Algorithm 8.3 will be equal to 1. For the choice of  $\theta$ , we used the choice made in [24]. If  $\theta$  is close to 1 then we maximize the use of the FP16 range and thereby reduce the chance of underflow (which in the worst case could make the matrix singular). On the other hand,  $\theta$  needs to be sufficiently less than 1 to allow headroom so subsequent computations do not overflow. We followed the study of [24] and used  $\theta = 0.1$ .

```

1. Perform  $\tilde{A} = RAC$ 
2. Let  $\beta$  be the maximum magnitude of any entry of  $\tilde{A}$ .
3. Let  $\mu = \theta x_{\max} / \beta$ 
4. scale  $A$  such that  $A^{(scaled)} = \mu \tilde{A}$ 

```

**Algorithm 8.3:** Two-sided scaling and scalar scaling.

We discuss the three scaling techniques described above in section 14 and show the benefits of each.

## 9 Multiple Right-Hand Side Optimizations

Many engineering applications require  $Ax = b$  to be solved for multiple right-hand sides (RHSs). For that, and in order to accommodate such applications, we designed and developed our Tensor Cores Accelerated Iterative Refinement Solver in such a way that allows for the efficient solution of multiple right-hand

sides. The TCAIRS consists of two main phases, the factorization phase and the refinement phase. The factorization phase performs an LU decomposition on the matrix  $A$  and is thus independent of the number of right-hand sides for which the linear system needs to be solved. As a result, this phase remains the same. The refinement phase is where the solution is refined to achieve the precision of the input data. In this work, we optimized and used the IR method described in Algorithm 4.1 for the solution of multiple right hand side. When multiple right-hand sides have to be solved, we need to adapt the computation of the residual to guarantee that all the solution vectors are accurate. This can be achieved if the error norm used for convergence checking is taken to be the maximal norm of all residuals (e.g., the maximal norm of each residual corresponding to each vector). This way, once the convergence is achieved, we guarantee that all solution vectors are accurate to the tolerance requested. However, since the refinement phase is memory bound, performing the refinement vector by vector is too expensive and will lead to a quick drop in the performance. Instead, we replace the matrix-vector operation `Xgemv` when computing the residual of  $b - Ax_i$  at iteration  $i$  with a matrix-matrix product `Xgemm` to compute all the residuals at once. We also replace the computation of an approximate solution using the LU factorization with the computation of all approximate solutions at once (e.g., replacing the two triangular solves over one vector using `Xtrsv` routines with two triangular solves for many vectors at once using the `trsm` routine). Moreover, to avoid computing the maximal norm of each residual, and then take the maximum of all of them, we only compute the norm of the first vector, and check if it satisfies the convergence error tolerance; if not, then we do not compute the remaining residual norm, but rather we step to the next iteration of the refinement loop, and keep doing the same process until that norm (norm of the first residual vector) satisfies the convergence error tolerance. It is at this point that we move to computing the next residual norm and follow the same process until all norms satisfy the convergence tolerance. To summarize, as soon as a residual norm does not satisfy the tolerance, we skip the norm computation of the remaining residuals and perform a new refinement iteration. This technique reduces the amount of memory-bound operations to be performed when multiple right-hand sides are requested.

## 10 Sensitivity of Performance to FP64 Compute Throughput

Because of the way the mixed-precision IR solver solves a problem—first using reduced and mixed-precision, and then by refining the solution to attain the FP64 accuracy with a memory bandwidth bound algorithm—it highlights the fact that it could be a good candidate to solve an FP64 problem with good performance on a GPU with Tensor Core capabilities, high memory bandwidth and not necessarily high FP64 throughput. This is theoretically demonstrated in detail in section 12 through the development of a model of the solver. We also then conducted such experiment by using the NVIDIA Turing GPU which is almost identical to Volta GPU except for FP64 compute throughput. This was chosen as a surrogate method to throttling only the FP64 compute throughput on Volta GPUs. We will show in section 15 (Figure 15.5) that our TCAIRS exhibits on a Quadro RTX8000 (Turing TU102 GPU) the same performance as on a high FP64 compute throughput card such as the Quadro GV100 (Volta GV100 GPU).

However, we should mention that applications in scientific computing typically do not only rely on solving a dense linear system  $Ax = b$ , and even when they need it, this is typically one component of the whole application. Generally, domain-specific algorithms, and other linear algebra functionality such as fast Fourier transforms (FFTs), eigenvalues, singular values, least squares, symmetric  $LDL^T$  decomposition, and matrix-matrix products are commonly used in applications. Thus to use GPU cards with low FP64 compute throughput for HPC scientific applications we will need many other mixed-precision numerical algorithm improvements.

## 11 Energy Efficient Implementation

Power efficiency in HPC is increasingly becoming a concern. Over the last few decades, the improving performance of HPC systems has come at the cost of increased electrical power consumption. The main concerns are increased power bills, e.g., going beyond affordable budgets, and increased impact on the environment. To help mitigate the power constraints in modern and future HPC systems, different approaches have been investigated. Exploiting both the hardware features and algorithms is an effective

solution to achieve power efficiency and to address those energy constraints. In this work we redesigned the solvers, which are typically the most time-consuming kernels in HPC applications, to provide energy-efficient alternatives. While most of the energy efficiency approaches aim to reduce the consumption with a minimal performance penalty, our TCAIRS improves both the performance and the energy efficiency. Indeed, we show below that by efficiently using the Tensor Core hardware and mixed precision, compared with highly optimized linear system solvers, our solver delivers the same accuracy solution and with more than 80% reduction in the energy consumption.

Also, our TCAIRS is a GPU-only implementation, which means it does not use the CPU for any computation. Thus, the CPU is idle, which also adds an advantage to the efficiency that such a solver can provide, in particular because CPUs require significant power consumption while providing significantly less performance. In other words, if CPUs contribute to the computation, their rate of FLOPs/Watt is very low, and thus we are going to observe a decrease in the energy efficiency of the solver. We will show in section 16 that our solver can effectively reach 94 gigaFLOP/s per Watt for FP64 real data and 126 gigaFLOP/s per Watt for FP64 complex data problems. These results significantly improve and extend previous work based on hybrid algorithms that use both CPUs and GPUs [16].

## 12 Performance Analysis

Mixed-precision methods derive their performance from the higher performance of lower-precision arithmetic. The theoretical peaks for the main precisions are shown in the last column of Table 1.1. In practice, the achievable  $X_{\text{gemm}}$  performance is less, but maintains a similar trend. Figure 12.1 quantifies the currently achievable performance for the NVIDIA GV100 GPU. In particular, we see that the PCIe GV100 has a practical peak of 6.8 teraFLOP/s in FP64, 14 teraFLOP/s in FP32, 28 teraFLOP/s in FP16, and a remarkable 85 teraFLOP/s in FP16-TC (Tensor Cores). However, the performance of the LU factorization relies mostly on the performance of the rank- $k$   $X_{\text{gemm}}$  updates, where the blocking size  $k$  is typically 256. These updates occur during each step of the LU algorithm, as given in Algorithm 4.2. Their performances are shown in Figure 12.1a (the dashed lines) for the four available precisions (FP64, FP32, FP16, and FP16-TC). Note that the rank- $k$   $h_{\text{gemm-TC}}$  achieves about 35 teraFLOP/s, compared to about 25 teraFLOP/s for the rank- $k$   $h_{\text{gemm}}$ , 13 teraFLOP/s for the rank- $k$   $s_{\text{gemm}}$ , and around 6 teraFLOP/s for the rank- $k$   $d_{\text{gemm}}$ . We note that in our work we always use the FP16-TC ( $h_{\text{gemm-TC}}$ ) and not the standard fused multiply-add (FMA) FP16  $h_{\text{gemm}}$ .

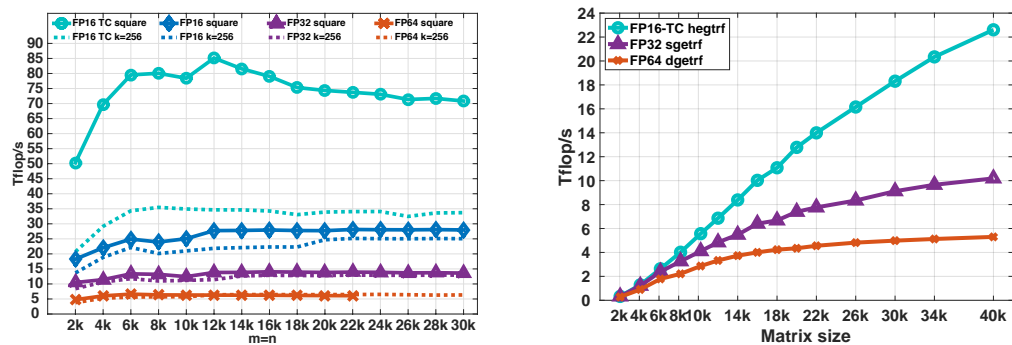
The multi-precision LU factorization from Algorithm 5.1, with  $\mathbf{u}^{\text{fh}}$  equal to FP32 and  $\mathbf{u}^{\text{fl}}$  equal to FP16-TC, was implemented and tuned for current hardware. Figure 12.1b shows the performance for the FP64, FP32, and FP16-TC precisions on GV100 GPUs using the cuSolver library [35]. As expected, the LU implementation follows roughly the same trend as the  $h_{\text{gemm-TC}}$  kernel for large  $n$ , showing that the implementation is well optimized and able to attain its theoretical upper bound. The  $h_{\text{getrf-TC}}$  solver achieves a speedup of up to  $4.5\times$  over  $d_{\text{getrf}}$  and a  $2.4\times$  speedup over  $s_{\text{getrf}}$ . The speedups further increase for matrix orders larger than 40,000.

Having the achievable performance  $P$  of the LU factorization, we can derive a theoretical performance analysis of the mixed-precision (MP) algorithms. This analysis is needed in order to understand and predict the cases where iterative refinement can be used. In particular, the interest is in cases where the solution can be reached faster than the reference (e.g., the FP64  $d_{\text{gesv}}$  routine). We recall that the iterative refinement solver performs an LU factorization in low precision, followed by a refinement loop based on either the classical IR, the IRGM, or the GM refinement method to improve the solution to  $\epsilon_{\text{FP64}}$ . Thus we can model performance for the real case by

$$\text{time for FP64} = \frac{2n^3}{3P_{d_{\text{getrf}}}} + \frac{2n^2}{P_{d_{\text{trsv}}}}, \quad (12.1)$$

$$\text{time for MP} = \frac{2n^3}{3P_{h_{\text{getrf-TC}}}} + k \left( \frac{2n^2}{P_{d_{\text{gemv}}}} + \frac{2n^2}{P_{s_{\text{trsv}}}} + \xi \right), \quad (12.2)$$

where  $P$  is the performance of the corresponding routine,  $k$  is the number of iterations to achieve the FP64 accuracy (including the inner GMRES iterations in the case of the IRGM solver), and  $\xi$  is other



(a) Performance of  $X_{gemm}$  for square matrices and for the rank- $k$  update used in the LU factorization ( $X_{getrf}$ ). (b) Performance of  $X_{getrf}$  for different precisions.

Figure 12.1: Performance for FP64, FP32, and FP16-TC routines on an NVIDIA V100 GPU.




Figure 12.2: Expected speedup of the iterative refinement solver over  $d_{gesv}$  as function of the number of iterations and the matrix size.

work required by the iterative refinement such as norm computation, residual calculation, pivoting, and synchronizations. Experiments show that  $\xi$  is negligible compared with the cost of the  $d_{gemv}$  and the  $strsv$ .

Based on the LU performance results provided in Figure 12.1b and on the benchmark of the  $d_{gemv}$  and  $strsv$  routine, we illustrate in Figure 12.2 the expected speedup of the iterative refinement solver using either the FP32 or the FP16-TC precision ( $d_{sagesv}$  and  $dh_{gesv-TC}$  respectively) as a function of the number of iterations (see also [17] for more details). In the latter case the expected speedup is up to a factor 4. We also note that the study for complex arithmetic is similar and the expected performance follows the same ratio, and for that we omit to redo the same model.

### 13 Experimental setup

Our experiments were performed on a system with one 10-core Intel(R) Xeon(R) E5-2650 v3 CPU running at 2.30 GHz equipped with an NVIDIA Quadro GV100 PCIe GPU and an NVIDIA Quadro RTX 8000 PCIe GPU.

To study the proposed methods and to highlight their practical use, we performed a large set of experiments on 21 types of matrices, with each type featuring different properties that represent a wide

Table 13.1: Description of the test matrices, where  $\text{cond}$  is  $\kappa_2(A)$ .

Type		Description
0	–	Random numbers with diagonal modified to be dominant
1	$\lambda > 0$	Random $\sigma$ in $[\frac{1}{\text{cond}}, 1]$ such that their logarithms are uniformly distributed
2	–	Random $\sigma$ in $[\frac{1}{\text{cond}}, 1]$ such that their logarithms are uniformly distributed
3	$\lambda > 0$	Clustered $\sigma$ : $\sigma = [1, \dots, 1, \frac{1}{\text{cond}}]$
4	–	Clustered $\sigma$ : $\sigma = [1, \dots, 1, \frac{1}{\text{cond}}]$
5	$\lambda > 0$	Arithmetically distributed $\sigma$ : $\sigma_i = 1 - (\frac{i-1}{n-1})(1 - \frac{1}{\text{cond}})$ , $i = 1, \dots, n$
6	–	Arithmetically distributed $\sigma$ : $\sigma_i = 1 - (\frac{i-1}{n-1})(1 - \frac{1}{\text{cond}})$ , $i = 1, \dots, n$
7	$\lambda > 0$	Geometrically distributed $\sigma$ : $\sigma_i = \text{cond}^{-\frac{i-1}{n-1}}$ , $i = 1, \dots, n$ ,
8	–	Geometrically distributed $\sigma$ : $\sigma_i = \text{cond}^{-\frac{i-1}{n-1}}$ , $i = 1, \dots, n$ ,

range of real life problems. We found that we could classify the 21 types of matrices using 6 representative cases that we present in Table 13.1. The matrices of singular vectors are random orthogonal matrices from the Haar distribution [44]. In the table,  $\lambda > 0$  denotes that the matrices have real, positive eigenvalues, and they are symmetric positive definite.

We also performed an extensive study on sparse matrices from the SuiteSparse Matrix Collection<sup>1</sup> [10].

## 14 Numerical Behavior

We first study the numerical behavior of our Tensor Cores Accelerated Iterative Refinement Solver using the three proposed refinement techniques, and show the convergence history of each technique for the different types of matrices. Then we discuss the benefit of using the Krylov GMRES solver. We also show the difference between using the IRGM and the GM technique. We then study the effect of scaling on the convergence of the solver for matrices from Table 13.1 as well as for the matrices from SuiteSparse.

In all our experiments,  $\mathbf{u}^{\text{tr}} = \mathbf{u}^{\text{fw}}$ .

### (a) Convergence Behavior

This study aims to provide an analysis of each method's sensitivity relative to a wide range of matrices representable for different scientific applications. In addition, our goal is to provide insight into the performance expected from the iterative refinement methods. For example, if an iterative refinement method requires a large number of iterations to achieve FP64 solution accuracy for a certain matrix type, then we can expect that its performance will degrade relative to the standard `dgesv`, and it may be even slower (see Section 12 for the expected performance as a function of the number of iterations). We note that the number of iterations that we report is the total number of iterations—including the inner GMRES iterations in the case of the IRGM solver. This means that the number of iterations reported is a precise indicator of the time spent in the refinement loop. The convergence criterion used in our experiments is the same as the one used in the state-of-the-art LAPACK iterative refinement solver `dsgesv`, namely  $\|Ax - b\| / (\|A\| \|x\|) < \epsilon^w \sqrt{n}$ , where  $\epsilon^w$  is the machine precision at the working precision (e.g.,  $\epsilon^{\text{FP64}}$  in our experiments) and  $n$  is the size of the matrix.

In the figures below, we plot the relative residual  $\|Ax - b\| / (\|A\| \|x\|)$

- at each iteration for the FP16-TC variant of the iterative refinement solver using the three proposed iterative refinement algorithms GM, (green) IR (yellow), and IRGM (blue),
- at each iteration for the FP32 variant using the GM algorithm (purple),
- of the final solution at the user level for the standard FP64 solver (horizontal orange line),

<sup>1</sup>Formerly known as the University of Florida Sparse Matrix Collection.



- of the final solution at the user level for the iterative refinement solver with the FP16-TC as lower precision (horizontal green, yellow and blues lines for the GM, IR and IRGM refinement methods respectively).
- of the final solution at the user level for the iterative refinement solver with the FP32 as lower precision (horizontal purple line),

The *no\_cvg* text in the figures mean that the corresponding refinement solver did not converge after 200 iterations.

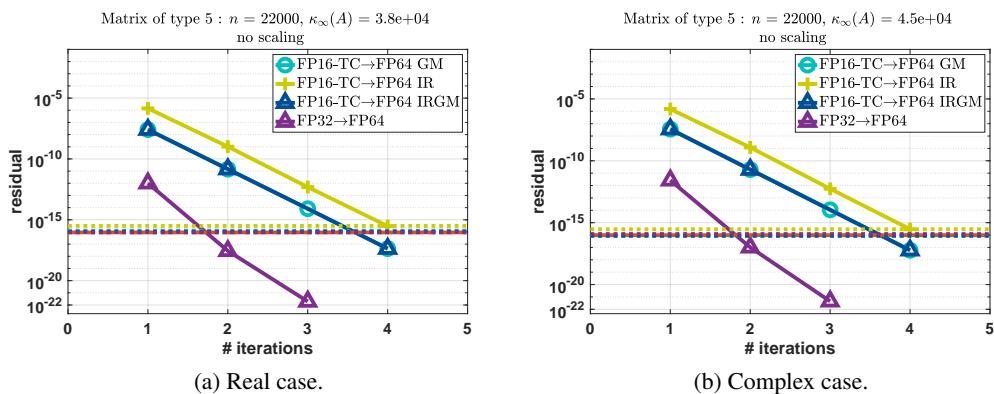


Figure 14.1: Convergence history of the proposed iterative refinement solver having either the FP32 as lower precision (using the GM refinement method) or the FP16-TC as lower precision ( using the GM, IR, and IRGM refinement methods). Matrix of size 22,000 × 22,000. Matrix of type 5: positive eigenvalues and arithmetic distribution of its singular values  $\sigma_i = 1 - \left(\frac{i-1}{n-1}\right)\left(1 - \frac{1}{\text{cond}}\right)$ . Similar behavior has been observed for matrices of types 0, 1, 3, and 7 within ± 1 iteration.

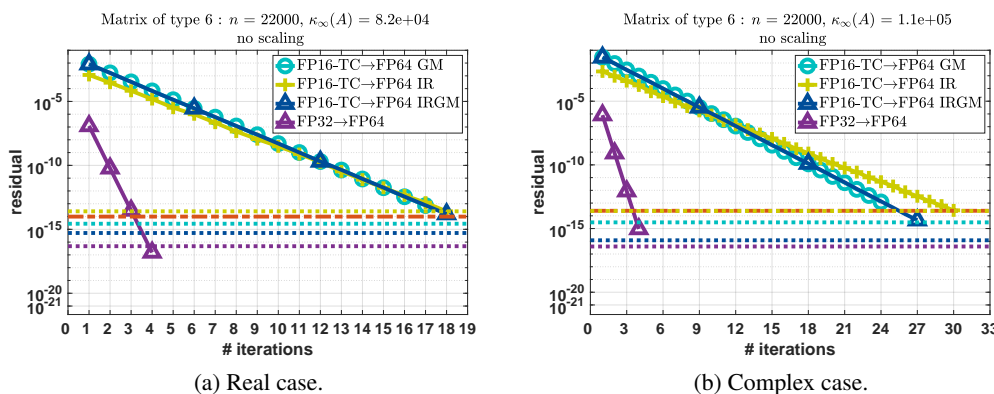


Figure 14.2: Convergence history of the proposed iterative refinement solver having either the FP32 as lower precision (using the GM refinement method) or the FP16-TC as lower precision ( using the GM, IR, and IRGM refinement methods). Matrix of size 22,000 × 22,000. Matrix of type 6: arithmetic distribution of its singular values  $\sigma_i = 1 - \left(\frac{i-1}{n-1}\right)\left(1 - \frac{1}{\text{cond}}\right)$ . Similar behavior has been observed for matrices of types 2, 4, and 8 within ± 2 iterations.



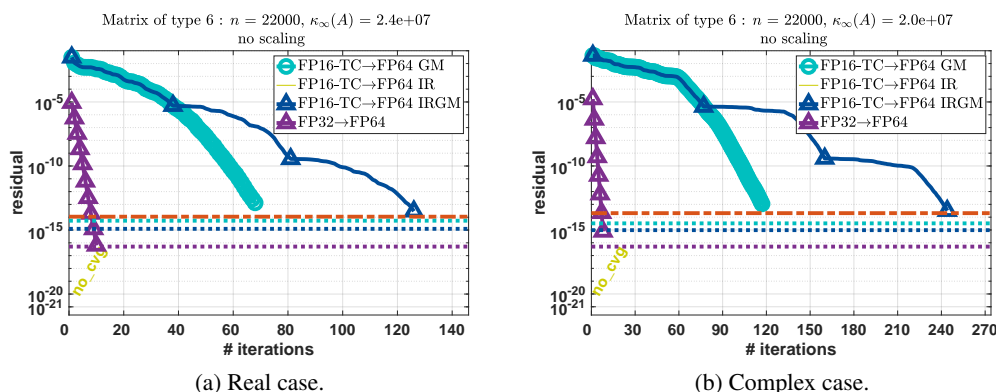


Figure 14.3: Convergence history of the proposed iterative refinement solver having either the FP32 as lower precision (using the GM refinement method) or the FP16-TC as lower precision (using the GM, IR, and IRGM refinement methods). Matrix of size 22,000 × 22,000. Matrix of type 6: arithmetic distribution of its singular values  $\sigma_i = 1 - (\frac{i-1}{n-1})(1 - \frac{1}{\text{cond}})$ . The matrix is generated with a large condition number to show the difference between the three proposed iterative refinement algorithms as well as to display the effect of the condition number on the convergence rate.

Figures 14.1 and 14.2 show the convergence history given by the relative residual at each step of the refinement loop for the six proposed iterative refinement solvers (the two precision implementations—FP16-TC and FP32—each using one of the three refinement methods GM, IR or IRGM). We note that when we use the FP32 as lower precision, the three refinement methods (i.e., GM, IR, and IRGM) behave exactly the same. Thus, to make graphs clear, we only show the GM implementations and denote them by FP32→FP64. The graphs are labeled as FPXX→FP64 YY, where “XX” corresponds to the lowest precision used during the LU factorization (FP16-TC or FP32), and “YY” represents the iterative refinement technique (GM, IR, or IRGM) used to attain the FP64 solution accuracy. In addition, in order to also compare the quality of the solution, we draw three horizontal lines that correspond to the residual of the solution at the output of our iterative refinement solver as well as the one at the output of the standard FP64 precision (dgesv and zgesv for real and complex cases, respectively).

In Figure 14.1, we display the convergence history of a matrix of type 5, but as mentioned in the caption, we observed the same behavior for the matrices of types 0, 1, 3, and 7 within ± 1 iteration difference. Here we can see that the TCAIRS using the FP32 as lowest precision converges within 3 iterations. Furthermore, we observe that the TCAIRS using the FP16-TC as lowest precision converges within 4 iterations for all refinement method variants (GM, IR, and IRGM). For these type of matrices, since the number of iterations is small, we can expect a large speedup over the FP64 routine. We believe that the FP32 solver will achieve a 2× speedup and the FP16-TC one will achieve about 4× speedup while delivering a solution at FP64 accuracy. More details about the performance are provided in the next section.

Figure 14.2 represents a matrix of type 6, which is very similar to type 5 in the sense that they both have the same singular values. However, their eigenvalues differ. We also mention that the same convergence trend has been observed for other matrices of types 2, 4, and 8. Thus, we omit the graphs for all of them. The matrices generated with types 2, 4, 6, and 8 characteristics are more difficult than the ones with types 0, 1, 3, 5, and 7. The convergence of FP16-TC requires more iterations, namely ≈ 17 iterations. Interestingly, the FP16-TC solver is doing well for this type of harder problem and is able to bring the solution down to the FP64 accuracy within an acceptable number of iterations. As a result, we could expect that the FP16-TC will provide about 3–4× speedup over the FP64 counterpart solver while delivering a solution at FP64 accuracy.

We mention that we performed some experiments with a basic FP16 implementation that does not use the Tensor Cores and we found that such implementation requires many more iterations to converge for this type of matrices. It also fails for other problems like the one in Figure 14.3 and many of the sparse problem

shown below. This is because the accumulation in the FP16-TC compared with the basic FP16 is done in FP32 arithmetic and thus produces a better result than the pure FP16. In addition we should not forget the fact that a basic FP16 (e.g., using the standard fused multiply-add (FMA)) will not exhibit the same high speed as the one with Tensor Cores. Such an algorithm and implementation will therefore be substantially slower, so we do not discuss it in this paper.

Figure 14.3 shows a substantially more difficult matrix. It is a matrix of type 5 where we increased the condition number to be within the FP32 range rather than the FP16-TC range. With such a configuration, even the FP32 will have trouble converging within 3 iterations, and will require about 10 iterations. For this type of matrix, and such a high condition number, the FP16-TC→FP64 IR variant is not able to converge. The basic IR method solving  $Ac_i = r_i$  directly using the computed LU factors and the solution is not accurate enough to allow the IR method to converge. However, both the IRGM and the GM refinement techniques (FP16-TC→FP64 IRGM, and FP16-TC→FP64 GM) are able to achieve the FP64 solution accuracy. This reveals the sensitivity of the IR variant and highlights the importance of using the preconditioned GMRES method. What happens here is that the GMRES method delivers a more stable solution to the refinement process (because the GMRES fixes the perturbed solution, thanks to the Krylov basis) which allows both the GM and the IRGM method to converge. We also note that the diamond marker in the green line—the curve that represents FP16-TC→FP64 IRGM—shows the number of outer iterations (outer refinement loop) in the IRGM solver. In our experiment, the GMRES tolerance of the IRGM method is set to  $10^{-4}$ . Thus, at each outer iteration of the refinement loop of the IRGM method the solution of  $Ac = r$  is correct to  $10^{-4}$  accuracy. We can see that the number of outer iterations is about 4, which is consistent with the theory of [7,8]. More details about using GMRES inside the refinement process can be found in [7,8].

Now, when comparing the two successful refinement methods in Figure 14.3 (i.e., GM and IRGM), we note that the GM method converges faster than the IRGM one. The GM method uses preconditioned GMRES to solve the whole linear system  $Ax = b$ , using the  $L$  and  $U$  factors as a preconditioner instead of using the preconditioned GMRES to solve the correction equation  $Ac = r$ , as in Section 6. In the IRGM method GMRES will stop at  $10^{-4}$ , then the outer refinement loop will compute a new residual, and then another GMRES is called to solve the new correction equation. Thus, the IRGM method can be viewed as a restarted GMRES—but instead of performing  $m$  iterations before each restart, it performs any number of iterations until it reaches the tolerance set for the GMRES (i.e.,  $10^{-4}$  in this study), and only then restarts. This is confirmed by the shape of the IRGM convergence history, and we can see that after each outer iteration (diamond in the green line) the GMRES starts constructing a new Krylov basis to solve the new correction equation  $Ac_i = r_i$ . This hard case example highlights the benefit of using a full GMRES to solve the system (e.g., the GM method) where the GMRES benefits from the entire previously constructed basis to compute the new direction for each new solution.

To make this study more general and cover many real life cases, we conducted the same study for sparse matrices arising from different engineering areas from the SuiteSparse Matrix Collection. The names, size, and condition numbers of these matrices are given in Table 15.1. We illustrate the convergence history results in Table 14.1 and Table 14.2 (the "no scaling" column). These matrices have condition numbers varying from  $10^1$  to  $10^{10}$ . For the iterative refinement solver using the FP32 as lower precision (Table 14.1), we mention that it converges for all of these cases, except for the problem "ramage02", where we also found that the scaling technique can be the remedy. Interestingly, we notice that the Tensor Cores Accelerated Iterative Refinement Solver using the FP16-TC precision as the lowest precision (Table 14.2) is able to converge for most of these problems and provide a solution similar to the one provided by the FP64 routine within a small number of iterations. These results highlight the benefit that the TCAIRS can bring to such applications, and one can expect a large speedup for these problems. We were surprised by the wide range of problems that the FP16-TC iterative refinement solver (TCAIRS) can solve. These results motivated us to look into further details and analyze why for some of these problems the FP16-TC variant had issues converging. We found that if we apply scaling to the matrix we can make the FP16-TC variant cope and solve these problems as well. This will be discussed below. We also mention that both the GM and the IRGM methods converged better than the standard IR method. Also, we notice the robustness and better convergence behavior that the GM refinement method showed compared to its IRGM counterpart. The GM method always converged similarly to or faster than the IRGM one. For hard cases, such as problem

"*sme3Db*" and "*ramage02*", we can observe the restarted shape of the GMRES of the IRGM method while the GM method illustrated around 40% less iterations.

**Lesson:** For the matrices considered, the TCAIRS using the FP16-TC precision is the most powerful method in terms of performance, and it is able to cope with and solve many problems representing many engineering areas within an acceptable number of iterations. The FP32 refinement variants show a more consistent behavior of small numbers of iterations regardless of the matrix types, except for minor cases. However, as shown in the analysis and in the performance sections, the maximum speedup that the FP32 variant can provide is less than  $2\times$  while the variant using FP16-TC can achieve  $4\times$ . These results suggest the surprising effectiveness and robustness of the FP16-TC arithmetic, showing it might be robust enough for use in HPC linear system solvers.

## (b) Scaling Techniques

As discussed above in Section 8, scaling has been studied and used in many direct and iterative full-precision solvers (e.g., a fixed-precision solver such as the FP64 solver) as a remedy to provide an accurate and stable solution for matrices with elements of widely varying magnitudes. In our work, as in the work of [24], we study scaling in the context of the iterative refinement solver and analyze its benefit. In Section 8 we proposed three scaling techniques and we denoted them by scalar scaling, diagonal scaling, and diagonal+scalar scaling. In this section we comment on the experimental results of these techniques. We first study these scaling approaches for the synthetic matrices described in Table 13.1 and then we comment on their benefit for the sparse problems from the SuiteSparse Matrix Collection.

**Synthetic matrices.** For the synthetic matrices of Table 13.1, we mention that only the scalar scaling need be used. The diagonal scaling technique proposed above computes a row and a column scaling intended to equilibrate the matrix  $A$  such that the largest element in each row and column of  $A$  has absolute value 1. Since the entries of the synthetic matrices are generated within a distribution between the interval  $[-1, 1]$ , there will be no row or column scaling to be found. As a result, the diagonal scaling will not change or scale the matrix  $A$ . For the same reason, the diagonal+scalar scaling is equivalent to just scalar scaling. Hence we illustrate only the effect of the scalar scaling for the synthetic matrices of Table 13.1. We choose the same matrices shown in Figure 14.1 and Figure 14.2 and perform the scalar scaling to the matrix  $A$  before we solve it using our iterative refinement solver. We illustrate the convergence history in Figure 14.4. The graph in Figure 14.4a is the convergence history of a matrix of type 5 without scaling, while the one on the right in Figure 14.4b corresponds to the convergence history of the same matrix but with the scalar scaling performed. As we can see, the scalar scaling does not provide any benefit here. Already for such matrices the iterative refinement solver converged very quickly and did not suffer from any perturbation. On the graph in Figure 14.4c, we represent the convergence history for a matrix of type 6 without any scaling and we depict the convergence history of the scaled matrix in Figure 14.4d. Here we can see that the scalar scaling cannot improve the convergence of the three refinement methods (GM, IR, and IRGM) of the FP16-TC variant, and, rather, it affects the numerics in a negative manner leading to a non-convergence. The reason for this behavior is that the scalar scaling increases the magnitude of the entries of  $A$  with the aim of minimizing the underflow—but since the entries of  $A$  become large, for matrices of this type (as well as matrices of types 2, 4, and 8), we believe that element growth greater than  $\theta^{-1} = 10$  causes elements of  $U$  to lie outside the FP16 range and then get rounded back onto the range, which reduces the quality of the LU factors, leading to non-convergence.

**Sparse matrices.** For the matrices from the SuiteSparse Matrix Collection we performed the three scaling techniques and illustrated the benefit of each in Table 14.1 and Table 14.2. The second block-column of the tables (from the left) corresponds to the scalar scaling, the third block-column corresponds to the diagonal scaling, and the fourth block-column (the right-most one) corresponds to applying the diagonal+scalar technique. As can be seen, the scalar scaling behaves similarly to the no-scaling for easy problems and it helped the convergence for harder cases. The diagonal scaling technique is more robust and showed benefit for most of the cases. We can see that for problems that were already easy to be solved by the iterative refinement solver without any pre-treatment, it did not slow down or affect the convergence. We also noticed how beneficial the diagonal scaling was for problems where the TCAIRS

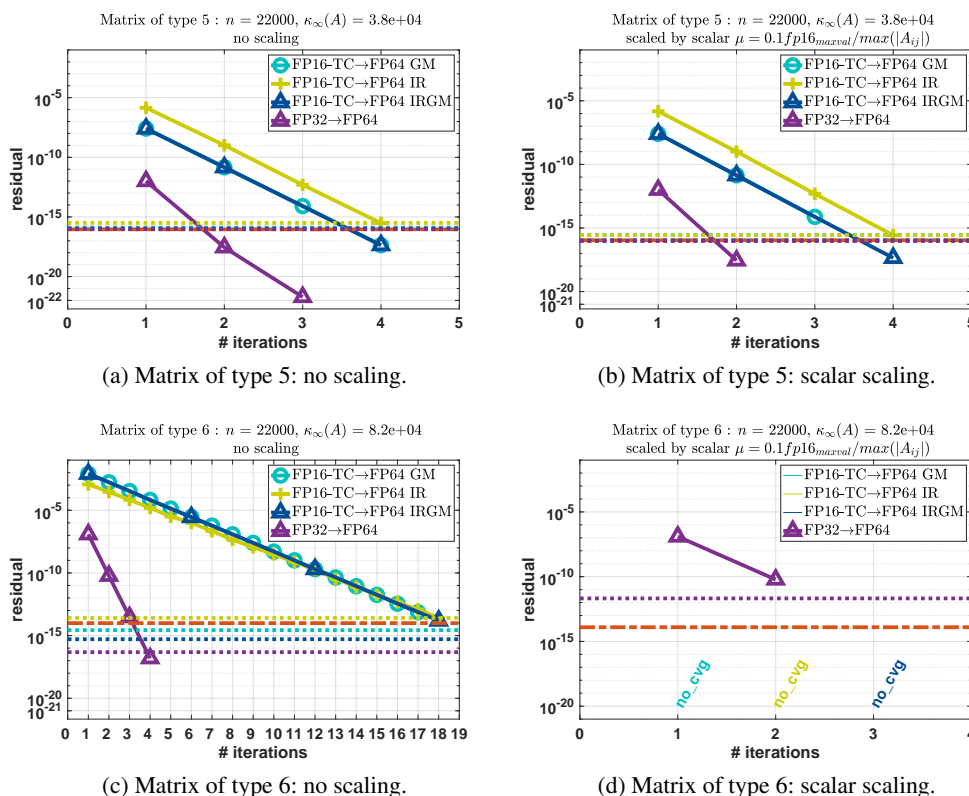


Figure 14.4: Effect of the different scaling techniques on the convergence of the iterative refinement solver. The example here, is for a matrix of type 5 (top) and type 6 (bottom) for real case of size of  $22,000 \times 22,000$ . We mention that complex cases behave similarly. Also, matrices of types 0, 1, 3, and 7 and type 2, 4, and 8 behave similarly to type 5 and type 6, respectively.

using the FP16-TC precision (e.g., Table 14.2) encountered trouble converging. The diagonal scaling helped the iterative refinement solver using the FP16-TC to converge on problems that the method was not able to solve without scaling. The scaling also helped decrease the number of iterations for problems where the TCAIRS without scaling converged but with a large number of iterations ("mixtank"). We also notice the decrease in the condition number of the matrices where diagonal scaling was applied. We show for each matrix, the condition number before and after scaling. It can be seen that when the iterative refinement solver using the FP16-TC precision did not converge (e.g., dashes in Table 14.2), the condition number of the matrices solved was very large  $> 10^8$ , and diagonal scaling helped decrease it, thereby enhancing the convergence of the FP16-TC iterative refinement solver. Similarly, for the problem "mixtank" for example, the FP16-TC iterative refinement solver converged within a medium number of iterations, but scaling accelerates convergence. The third strategy of scaling was to combine the two scaling techniques together. Our experiments showed that this strategy showed an unpredictable behavior. It helped some problems but also was worse than no scaling from some others.

**Lesson: Scaling techniques can enhance the convergence of the iterative refinement solver, in particular for problems with large a condition number. The best scaling technique is the equilibrium of the matrix which is denoted by the diagonal scaling. The scalar scaling might help the convergence for many of these cases, as long as the value of  $\theta$  is taken less than the reciprocal of the growth factor, thus  $\mu$  requires to be carefully chosen. The diagonal+scalar scaling showed that it behaves like scalar scaling for most of the cases, except for some cases, it was worse than both scaling strategies.**

Table 14.1: Number of iterations required by the proposed iterative refinement solver having the FP32 as lower precision (*dsgesv*) using the GM, IR, and IRGM refinement methods. Also, we show the effect of the different scaling techniques in each block-column.

Name	Size	$\kappa_{\infty}(A)$	no scaling			scalar scaling			diagonal scaling				diag+scalar scaling		
			IR	IRGM	GM	IR	IRGM	GM	$\kappa_{\infty}(A^*)$	IR	IRGM	GM	IR	IRGM	GM
crystk03	24696	2.09e+02	1	1	2	1	1	2	2.5e+01	1	1	2	1	1	2
crystm03	24696	8.24e+01	1	1	2	1	1	2	7.7e+00	1	1	2	1	1	2
TEM27623	27623	4.70e+03	1	1	1	1	1	2	2.1e+01	1	1	2	1	1	2
Goodwin054	32510	4.49e+06	2	2	3	2	2	3	2.2e+04	2	2	3	2	2	3
thermal1	17880	1.60e+03	1	1	3	1	1	2	8.0e+01	1	2	3	2	2	2
Zhao1	33861	4.45e+02	1	1	3	2	1	2	1.6e+01	2	2	2	2	2	2
nd6k	18000	3.82e+02	1	1	2	1	1	2	1.0e+02	1	1	2	1	1	2
nd12k	36000	3.51e+02	1	1	2	1	1	2	1.0e+02	1	1	2	1	1	2
epb1	14734	2.62e+04	2	2	4	2	2	3	4.4e+04	2	2	3	2	2	3
appu	14000	1.02e+04	2	2	3	2	2	2	5.8e+03	2	2	3	2	2	2
ns3Da	20414	3.79e+03	2	2	3	2	2	2	3.4e+03	2	2	3	2	2	2
mixtank new	29957	4.69e+11	2	2	4	2	2	3	2.1e+05	2	2	4	2	2	3
e40r0100	17281	2.23e+08	2	3	4	3	4	4	1.1e+04	2	2	3	2	2	2
bcsstk25	15439	1.02e+10	1	1	1	1	1	2	6.2e+02	1	1	2	1	1	2
bcsstk37	25503	2.27e+09	1	1	2	1	1	2	1.5e+02	1	1	2	1	1	2
raefsky3	21200	5.18e+11	1	1	3	1	1	2	1.2e+04	2	2	3	2	2	2
thread	29736	1.78e+08	1	1	1	1	1	2	1.3e+03	1	1	2	1	1	2
sme3Db	29067	1.05e+08	15	10	9	15	11	9	3.8e+07	16	11	12	16	11	9
mult dcop 01	25187	9.55e+11	3	2	6	3	4	4	1.5e+08	2	2	5	2	2	3
ramage02	16830	4.52e+08	-	60	39	-	58	38	4.5e+07	11	12	15	37	28	22

Table 14.2: Number of iterations required by the proposed iterative refinement solver having the FP16-TC as lower precision (*dhgesv-TC*) using the GM, IR, and IRGM refinement methods. Also, we show the effect of the different scaling techniques in each block-column.

Name	Size	$\kappa_{\infty}(A)$	no scaling			scalar scaling			diagonal scaling				diag+scalar scaling		
			IR	IRGM	GM	IR	IRGM	GM	$\kappa_{\infty}(A^*)$	IR	IRGM	GM	IR	IRGM	GM
crystk03	24696	2.09e+02	1	1	2	1	1	2	2.5e+01	1	1	2	1	1	2
crystm03	24696	8.24e+01	1	1	2	1	1	2	7.7e+00	1	1	2	1	1	2
TEM27623	27623	4.70e+03	1	1	1	1	1	2	2.1e+01	1	1	2	1	1	2
Goodwin054	32510	4.49e+06	8	5	6	8	5	5	2.2e+04	5	5	6	9	5	4
thermal1	17880	1.60e+03	3	4	4	3	4	3	8.0e+01	3	3	4	3	3	3
Zhao1	33861	4.45e+02	3	4	4	3	4	3	1.6e+01	3	4	4	3	4	3
nd6k	18000	3.82e+02	3	4	4	3	4	3	1.0e+02	2	2	3	2	2	2
nd12k	36000	3.51e+02	3	4	4	3	4	3	1.0e+02	1	1	2	1	1	2
epb1	14734	2.62e+04	3	4	4	3	4	3	4.4e+04	3	4	4	4	4	3
appu	14000	1.02e+04	4	4	5	3	4	3	5.8e+03	4	4	5	3	4	3
ns3Da	20414	3.79e+03	3	4	4	3	4	3	3.4e+03	3	4	4	4	-	-
mixtank new	29957	4.69e+11	-	36	31	13	11	10	2.1e+05	8	9	8	-	-	-
e40r0100	17281	2.23e+08	44	9	11	-	9	8	1.1e+04	5	6	6	-	-	-
bcsstk25	15439	1.02e+10	-	-	-	1	2	4	6.2e+02	3	4	4	3	4	3
bcsstk37	25503	2.27e+09	-	-	-	2	2	4	1.5e+02	3	4	4	3	4	3
raefsky3	21200	5.18e+11	-	-	-	4	4	11	1.2e+04	7	6	8	6	6	5
thread	29736	1.78e+08	-	-	-	2	2	3	1.3e+03	3	4	4	3	3	3
sme3Db	29067	1.05e+08	-	15	9	74	13	10	3.8e+07	-	20	11	-	18	9
mult dcop 01	25187	9.55e+11	-	-	-	-	108	62	1.5e+08	-	47	34	-	63	26
ramage02	16830	4.52e+08	-	-	-	-	-	-	4.5e+07	-	179	83	-	-	-

## 15 Performance

This section presents the performance results of our two proposed iterative refinement solvers with the FP32 or the FP16-TC as lower precision—*dhgesv-TC*, *dsgesv* and *zkgesv-TC*, *zcgsv* for real and complex cases, respectively—using the GM (GMRES) refinement method, and comparing it to the full precision reference *dgesv* and *zgesv* solvers, respectively. In all the performance figures below, we illustrate the performance in teraFLOP/s of three linear solvers for both real and complex cases:

- (1) the FP64 standard solver (*dgesv* and *zgesv* for real and complex, respectively), displayed in orange color with “×”,
- (2) the iterative refinement solver having the FP32 precision as lower precision (i.e.,  $u^f=FP32$ ) and using the GM refinement method (*dsgesv* and *zcgsv* for real and complex matrices, respectively) displayed in purple color with “△”,
- (3) the Tensor Cores Accelerated Iterative Refinement Solver using having the FP16-TC precision as lower precision (e.g.,  $u^f=FP16-TC$ ) and using the GM refinement method (*dhgesv-TC* and *zkgesv-TC* for real and complex matrices, respectively) displayed in cyan color with “o”.

We mention that the teraFLOP/s are computed using the same formula for all solvers and all plots,  $P = \frac{2n^3}{3t}$  where  $t$  is the total elapsed time of the computation, which means performance reflects the time to solution. Thus, a performance that is two times higher means the computation is twice as fast. We also note that the time of the iterative refinement solver includes all the conversion required and the iterations, meaning the solver takes a full precision  $A, x, b$  (e.g., FP64  $A, b$ ), and returns the solution  $x$  in full precision (FP64 in our experiments). The study was done for both real and complex matrices and for different matrix sizes and for different matrix types. For the two iterative refinement solvers, we also depict the required number of iterations to achieve the FP64 arithmetic solution. To make the figures fully informative of all the possible information that the reader might find interesting, we show in the right “y” axis of the graph the condition number  $\kappa_\infty(A)$ , corresponding to the gray dotted line.

**(a) Single Right-Hand Sides**

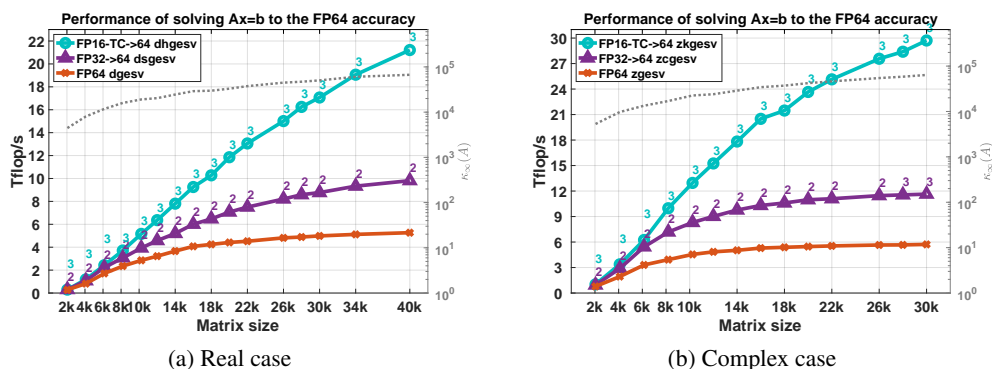


Figure 15.1: Performance in teraFLOP/s of the GM refinement algorithm for the FP32 and FP16-TC precisions studied for different matrix sizes on an NVIDIA GPU Volta GV100. Matrix of type 5: positive eigenvalues and arithmetic distribution of its singular values  $\sigma_i = 1 - \left(\frac{i-1}{n-1}\right)\left(1 - \frac{1}{\text{cond}}\right)$ . Similar behavior has been observed for matrices of types 0, 1, 3, and 7.

The performance results for the synthetic matrices of Table 13.1 are summarized in Figures 15.1-15.3. In Figure 15.1, the matrix is of type 5, and—as shown in Section 14—the two precisions variants (FP32 and FP16-TC) of the iterative refinement solver require two to four iterations to converge. Thus, one can expect that the iterative refinement solver will bring a large speedup compared to its full-precision counterpart solver. Since the number of iterations is small, we presume that the speedup ratio will be similar to the one observed in Figure 12.1b for the LU factorization. Our expectation is confirmed by the experimental results presented in Figure 15.1. The Tensor Cores Accelerated Iterative Refinement Solver using the FP16-TC as low precision (*dhgesv-TC* and *zkgesv-TC*) delivers a solution 4× and 5× faster than its FP64 counterpart *dgesv* and *zgesv*, respectively. Similarly, the iterative refinement solver variant using the FP32 as lower precision (e.g., *dsgesv* and *zcgsv* for real and complex data, respectively) shows a  $\approx 1.8\times$  speedup over their *dgesv* and *zgesv* FP64 precision counterpart. This example illustrates the



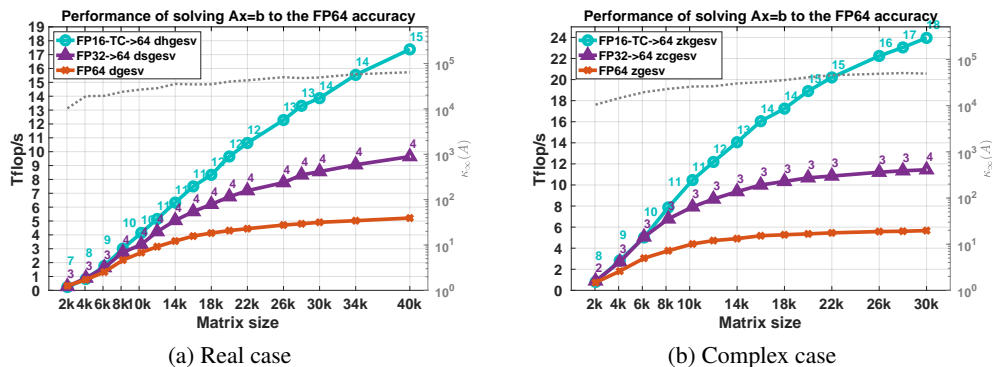


Figure 15.2: Performance in teraFLOP/s of the GM refinement algorithm for the FP32 and FP16-TC precisions studied for different matrix sizes on an NVIDIA GPU Volta GV100. Matrix of type 4: clustered singular values,  $\sigma_i=(1, \dots, 1, \frac{1}{\text{cond}})$ . Similar behavior has been observed for matrices of type 4.

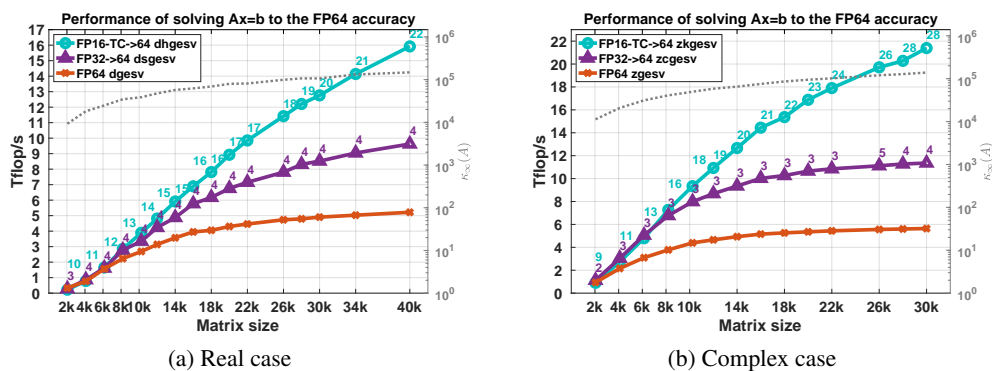


Figure 15.3: Performance in teraFLOP/s of the GM refinement algorithm for the FP32 and FP16-TC precisions studied for different matrix sizes on an NVIDIA GPU Volta GV100. Matrix of type 6: arithmetic distribution of its singular values  $\sigma_i = 1 - \left(\frac{i-1}{n-1}\right)\left(1 - \frac{1}{\text{cond}}\right)$ . Similar behavior has been observed for matrices of types 2 and 8.

importance of using the low FP16-TC precision in HPC. We note that our experiments showed that matrices of types 0, 1, 3, and 7 exhibited the same performance behavior as the one illustrated in Figure 15.1 (e.g., matrices of type 5). Thus, we omit 8 more graphs that are going to look the same. The performance results in Figure 15.3 for type 4 matrices are similar to those in Figure 15.1 for type 5 matrices. For the real case (Figure 15.2a), the iterative refinement solver dhgesv-TC and dsgesv outperform the dgesv by around 3.5× and 1.7×, respectively. For the complex case (Figure 15.2b), the iterative refinement solver zkgesv-TC and zcgesv outperform the dgesv by around 4.5×, and 1.8×, respectively.

In contrast to the matrices of type 5 shown in Figure 15.1, Figure 15.3 shows the performance and the number of iterations for matrices that have the same arithmetic distribution of the singular value, but their eigenvalues are not necessarily positive and can even be complex. This is considered a harder problem than the one shown in Figure 15.1. For the TCAIRS using the FP16-TC precision as lower precision (dhgesv-TC and zkgesv-TC), the number of iterations increases compared with Figure 15.1 (15–24 iterations vs. 3–5 iterations). We also see that the iteration count increases slightly with matrix size. Thus, one can expect the performance here to be slightly lower than in the one of Figure 15.1 (while still being about 3× faster). The dsgesv and the zcgesv behavior stays the same as the one shown in the previous

graphs and requires about 2–3 iterations, resulting in a  $1.7\times$  speedup over the `dgesv` and the `zgesv`, respectively. The observations made here are correlated with the numerical behavior shown in Section 14, Figure 14.2. It highlights the attractive effectiveness of the FP16–TC precision and reveals the interest of using it for HPC computing. Matrices of type 2 and type 8 have shown same performance behavior and thus we omit their graphs.

Table 15.1: Performance for real life matrices from the SuiteSparse Collection and from dense matrices arising from radar design simulations.

Name	Description	Size	$\kappa_{\infty}(A)$	dgesv	dsgesv			dhgesv-TC		
				FP64	FP32 → FP64	FP16-TC → FP64	# iter	time (s)	speedup	# iter
crystk03	Materials problem	24696	2.09e+02	1.95	2	1.16	1.68	2	0.69	2.82
crystm03	Materials problem	24696	8.24e+01	1.95	2	1.16	1.68	2	0.69	2.82
TEM27623	3D electromagnetic	27623	4.70e+03	2.68	2	1.54	1.74	2	0.87	3.08
mixtank new	CFD problem	29957	4.69e+11	3.56	4	2.07	1.71	8	1.16	3.06
e40r0100	2D/3D problem	17281	2.23e+08	0.81	3	0.53	1.52	6	0.39	2.07
sme3Db	Structural mechanics	29067	1.05e+08	3.33	12	2.09	1.59	11	1.13	2.94
Goodwin 054	Navier-Stokes	32510	4.49e+06	4.17	3	2.33	1.78	6	1.30	3.20
ns3Da	3D Navier Stokes	20414	3.79e+03	1.29	3	0.81	1.59	4	0.48	2.68
Thermal1	Thermal Problem	17880	1.60e+03	0.85	3	0.55	1.54	4	0.38	2.23
Zhao1	Electromagnetics	33861	4.45e+02	4.76	2	2.60	1.83	4	1.40	3.40
nd6k	2D/3D problem	18000	3.82e+02	0.86	2	0.55	1.56	3	0.38	2.26
nd12k	2D/3D problem	36000	3.51e+02	5.57	2	3.04	1.83	2	1.58	3.52
epb1	Thermal problem	14734	2.62e+04	0.53	3	0.36	1.47	4	0.27	1.96
appu	Graph problem	14000	1.02e+04	0.50	3	0.35	1.42	5	0.24	2.08
bcstsk25	Structural problem	15439	1.02e+10	0.59	2	0.39	1.51	4	0.29	2.03
bcstsk37	Structural problem	25503	2.27e+09	2.15	2	1.27	1.69	4	0.75	2.86
raefsky3	CFD problem	21200	5.18e+11	1.33	3	0.82	1.62	8	0.57	2.33
thread	Threaded connector	29736	1.78e+08	3.30	2	1.85	1.78	4	1.05	3.14
mult dcop 01	Circuit simulation	25187	9.55e+11	2.09	5	1.27	1.64	34	1.16	1.80
ramage02	Navier-Stokes	16830	4.52e+08	0.76	15	0.58	1.31	83	0.96	0.79
Poisson	2D Poisson problem	32000	2.1e+06	4.01	2	2.15	1.81	10	1.23	3.33
Vlasov	2D Vlasov problem	22000	8.3e+03	1.70	2	0.95	1.78	3	0.50	3.40

The goal of this paper is to show that the proposed iterative refinement solver can be of great benefit for a wide range of matrices with different characteristics. In practice, the real world matrices tend to be *easier* to deal with than our specially constructed synthetic ones. To illustrate this, we show results for real world matrices arising from different problems. We show in Table 15.1 the results from experiments obtained when running three linear solvers (the standard solver `dgesv` and our two iterative refinement



solvers using the FP32 and the FP16-TC precision, respectively) for different real world matrices from the SuiteSparse Matrix Collection.

As can be seen, the behavior of the proposed iterative refinement solver is similar to that of the synthetic matrices described above. For most of these problems, the iterative refinement solver using the FP16-TC as lower precision converges within fewer than 12 iterations, except for only 2 hard problems. For these hard problems, the iterative refinement solver using the FP32 as lower precision also had trouble converging within 2-3 iterations as usual. Overall, our proposed iterative refinement solver showed attractive speedup for these problems. By studying the results in the table we found that the proposed TCAIRS using the FP16-TC precision (e.g., the `dhgesv-TC` routine) can provide in average a speedup of about  $2.5\times$ – $3.5\times$  for a wide range of real world matrices and it shows about  $1.7\times$  speedup when using the FP32 as lower precision (e.g., the `dsgesv` routine). These results match the analysis performed in the numerical Section 14 and in the performance model Section 12.

**Lesson:** The speedups presented in Figures 15.1–15.3 and in Table 15.1 show that the `dhgesv-TC` (FP16-TC → FP64) routine can be used for most of the matrices types to provide speedups of about  $3\times$ – $4\times$ , and the `dsgesv` (FP32 → FP64) routine can be used for most of the matrix types to provide speedups of about  $1.7\times$ . The complex routines (`zkgesv-TC` and `zcgesv`) also showed similar behavior as the real ones (`dhgesv-TC` and `dsgesv`) but with a slightly higher speedup over the FP64 complex routine `zgesv`. This is due to the higher compute intensity of the complex data computations.

### (b) Multiple Right-Hand Sides

In this section we show the benefit of the Tensor Cores Accelerated Iterative Refinement Solver when solving multiple right-hand sides. As discussed above, the iterative refinement solver consists of two phases: the factorization phase, which uses lower precision to achieve a high performance, and a second phase to refine the obtained solution down to the accuracy of the FP64 precision. The refinement phase is memory bound and thus it is better to converge fast (see Figure 12.2 for the expected performance as function of the number of iterations). When multiple right-hand sides are needed, we cannot solve each right-hand side by itself because then we multiply the number of iterations by the number of RHSs. Thus, our strategy is to solve all the RHS together at once and to use Level 3 BLAS (`trsm` and `xgemm`) during the refinement in such a way that solving for one or more RHSs will take roughly similar time (with  $\approx 10\%$  difference). In Figure 15.4 we illustrate the performance obtained when solving 1 or 32 RHSs for a matrix of type 5.

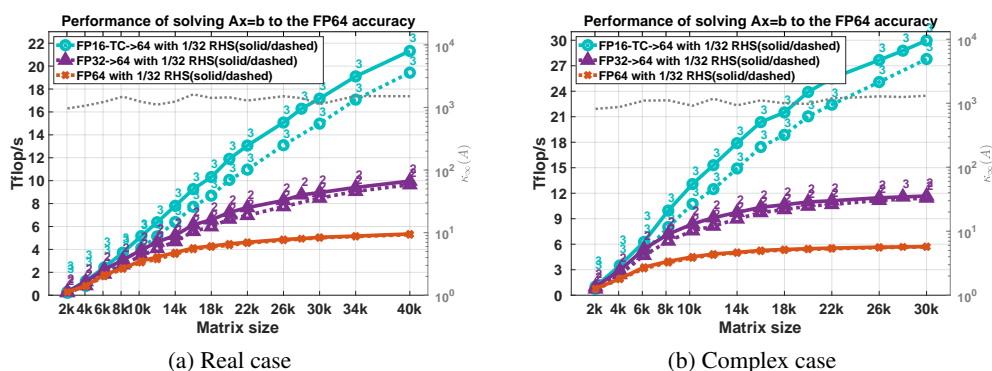


Figure 15.4: Performance of solving multiple right-hand sides using the three linear solvers for a matrix of type 3, for different matrix sizes on an NVIDIA GPU Volta GV100. Results are shown for 1 and 32 right-hand sides.

As observed in Figure 15.4 that solving 32 RHSs does not significantly delay the refinement process, and

thus the performance obtained is only less than 10% slower than the one obtained when solving 1 RHS, for both real and complex cases. This is thanks to our proposed technique that uses Level 3 BLAS operations and optimizes the residual checking process in such a way that minimises the amount of memory-bound operations. Interestingly, we can see that the number of iterations required to converge the 32 RHSs is similar to the one required to converge 1 RHS. These results give the TCAIRS an advantage even when multiple RHSs are to be solved.

### (c) Sensitivity to FP64 Compute Throughput

Based on the discussion in Section 10, we used our iterative refinement solver to solve the same problems studied in this paper on an NVIDIA Quadro RTX8000 (Turing TU102 GPU). As expected from the discussion above, we found that the performance of the TCAIRS is not sensitive to the FP64 compute throughput and maintains the same high performance behavior. Figure 15.5 depicts the performance obtained for a matrix of type 5 using the three solvers—the FP64 solver (*dgesv* or *zgesv* for real and complex cases, respectively), the iterative refinement solver with FP32 as lower precision (*dsgesv* or *zsgesv* for real and complex cases, respectively), and the iterative refinement solver using the FP16-TC precision as lower precision (*dhgesv-TC* or *zkgesv-TC* for real and complex cases, respectively)—on the GV100 and the RTX8000. We can easily see that the solver maintains the same effective FP64 performance on both GPU architectures despite their vastly different FP64 compute throughput (7 TFLOP/S on GV100 and 500 GFLOP/s on RTX8000). Here we show only one graph of performance, but all the other graphs show a similar performance trend.

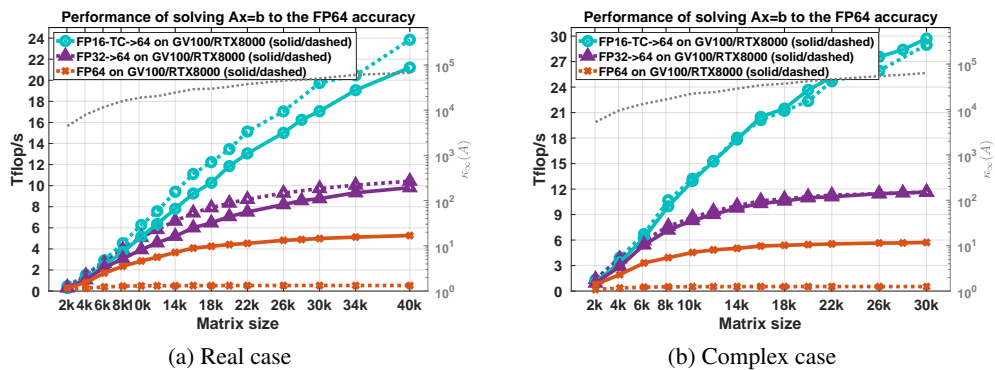


Figure 15.5: Performance comparison of the Tensor Cores Accelerated Iterative Refinement Solver on two different GPUs: NVIDIA Quadro GV100 and RTX8000 shows lack of sensitivity to FP64 compute throughput. We illustrate the iterative refinement algorithms using either the FP32 and FP16-TC precisions as lower precision. Matrix of type 5: positive eigenvalues and arithmetic distribution of its singular values  $\sigma_i = 1 - \left(\frac{i-1}{n-1}\right)\left(1 - \frac{1}{\text{cond}}\right)$ . Similar behavior has been observed for all other synthetic and real life matrices.

## 16 Energy Efficiency

Figure 16.1 shows the energy efficiency results in real (Left) and complex (Right) precision, respectively. We report the sum of the CPU, DRAM, and GPU power measurement. The power of the CPU (package+DRAM) is measured using the Performance Application Programming Interface (PAPI) [30], and the power of the GPU is measured using the NVIDIA Management Library (NVML) [36]. We note that the solver is GPU-only (i.e., does not use the CPU), but we still add the CPU's idle power consumption, which is about 20–30 Watts. The standard *dgesv* solver provides an energy efficiency of 21 gigaFLOP/s per Watt. The FP32 iterative refinement solver (*dsgesv*) doubles the energy efficiency, increasing it to 40

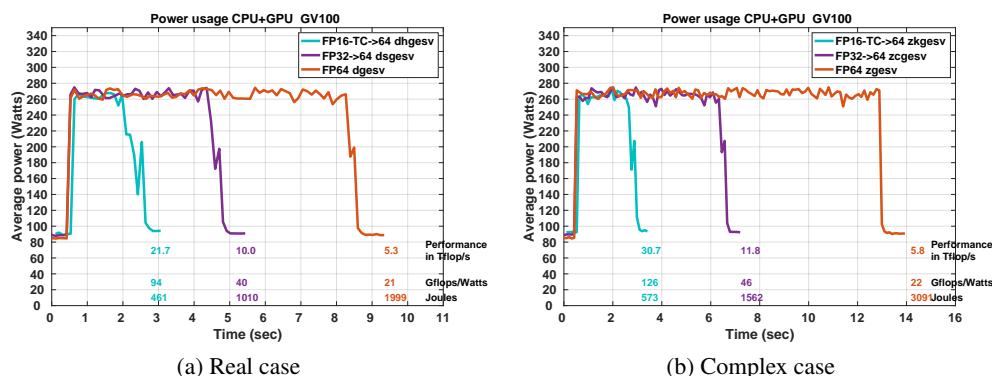


Figure 16.1: Power usage of the CPU and an NVIDIA GPU Volta GV100 for the Tensor Cores Accelerated Iterative Refinement Solver for the FP32 and FP16-TC precisions for a real and complex matrices of size  $n = 40000$  and  $n = 30000$ , respectively.

gigaFLOP/s per Watt. This follows our performance analysis described above, since *dsgesv* is about twice as fast and thus we can observe twice the energy efficiency using the *dsgesv* routine. The most pronounced result is achieved by the FP16-TC iterative refinement solver (*dhgesv-TC*). It achieves an unprecedented energy efficiency of 94 gigaFLOP/s per Watt, which is about  $4.5\times$  improvement over the standard *dgesv* solver. The complex case achieves 126 gigaFLOP/s per Watt, which is more than  $5\times$  improvement over the standard *zgesv* solver. These results demonstrate that the iterative refinement methods and half-precision arithmetic will be decisive in helping mitigate the power constraints in large-scale HPC systems.

The different phases of the computation can be identified on the power graph, as we now explain. The portion that consumes very high power (e.g., 260 Watts) corresponds to the factorization phase. 90% of the factorization phase consists of compute intensive operations and thus the power will be at the peak. We can also see here that this portion that correspond to the factorization in FP16 (cyan curve) is about 4 times faster than its corresponding one in FP64 (orange curve). Once the factorization is done, the code synchronize all streams, and performs some checking. This corresponds to the lowest quick spike down. Then, for the standard FP64 solver (e.g., *dgesv*), it performs the solve phase (one solve of  $LUx = b$ ) to compute the final solution  $x$  while for the iterative refinement solver, it performs the refinement phase (which correspond to as many solve and mat-vec as the number of iterations). The solve or refinement phase consists of memory bound operations and we know that such operations do not consumes the peak power but rather they consume lower power which correspond to the curves around 180 Watts. Thus we can easily illustrate the refinement phase in these graph by the portion that consumes around 180 Watts. We can see that the refinement portion for FP32 *dsgesv* is slightly shorter than the one for FP16 *dhgesv-TC*. This is normal since the *dsgesv* required about 3 iterations while the *dhgesv-TC* required about 7 iterations.

## 17 Conclusions and Future Directions

Designing and implementing numerical algorithms that efficiently exploit current highly parallel computer architectures is a challenge, especially if close to peak performance is to be achieved. Nevertheless, we have developed a new class of iterative refinement solvers and a number of computational techniques that allow us to solve fundamental  $Ax = b$  problems not just close to peak FP64 performance, but to get multiple times over the peak, by using the fast mixed-precision arithmetic Tensor Cores available in new GPU architectures—all while retaining the numerical stability of an FP64 solution. In particular, we showed that the new iterative refinement solver can accelerate the solution 4-to- $5\times$  and have  $5\times$  better energy efficiency on NVIDIA Volta GV100 GPUs. The complex case achieves an unprecedented energy efficiency of 126 gigaFLOP/s per Watt, which is more than  $5\times$  improvement over the standard FP64 solver. The new developments further improves previous efforts in this direction, including performance, energy efficiency,

and applicability to real problems. The solvers are now released in the vendor-optimized numerical library CUSOLVER [35] from the NVIDIA and the MAGMA [27] open-source numerical library.

The developments open up directions for future work, including further optimizations, development of a full set of mixed-precision factorizations, linear system solvers as well as eigensolvers and singular value decomposition (SVD), and release as open-source software through MAGMA [27]. See [23] for a mixed-precision algorithm for symmetric positive definite systems and [9], [23] for mixed-precision algorithms for the least square problem. Furthermore, the developments illustrate that mixed-precision techniques can be of great interest for linear solvers in many engineering areas. Such methods can also be easily ported to large-scale distributed or multi-GPU environments and supercomputers, where the speedups are expected to remain the same as for a single GPU.

Data Accessibility. This article has no additional data.

Authors' Contributions. All authors drafted and revised the manuscript. All authors read and approved the manuscript.

Competing Interests. The authors have no competing interests.

Funding. The work of AH and HB was supported by NVIDIA. The work of ST and JD was supported by the Exascale Computing Project (17-SC-20-SC), a joint project of the U.S. Department of Energy's Office of Science and National Nuclear Security Administration. The work of NJH was supported by Engineering and Physical Sciences Research Council grant EP/P020720/1 and the Royal Society.

## References

1. E. Anderson, Z. Bai, C. Bischof, L. S. Blackford, J. Demmel, Jack J. Dongarra, J. Du Croz, S. Hammarling, A. Greenbaum, A. McKenney, and D. Sorensen. *LAPACK Users' guide (third ed.)*. SIAM, Philadelphia, PA, USA, 1999.
2. M. Arioli, J. W. Demmel, and I. S. Duff. Solving sparse linear systems with sparse backward error. *SIAM J. Matrix Anal. Appl.*, 10(2):165–190, 1989.
3. Marc Baboulin, Alfredo Buttari, Jack Dongarra, Jakub Kurzak, Julie Langou, Julien Langou, Piotr Luszczek, and Stanimire Tomov. Accelerating scientific computations with mixed precision algorithms. *Computer Physics Communications*, 180(12):2526–2533, 2009.
4. Pierre Blanchard, Nicholas J. Higham, Florent Lopez, Theo Mary, and Srikara Pranesh. Mixed precision block fused multiply-add: Error analysis and application to GPU tensor cores. MIMS EPrint 2019.18, Manchester Institute for Mathematical Sciences, The University of Manchester, UK, September 2019.
5. Ian Buck. World's fastest supercomputer triples its performance record. <https://blogs.nvidia.com/blog/2019/06/17/hpc-ai-performance-record-summit/>, June 2019. Accessed June 17, 2019.
6. Alfredo Buttari, Jack Dongarra, Julie Langou, Julien Langou, Piotr Luszczek, and Jakub Kurzak. Mixed precision iterative refinement techniques for the solution of dense linear systems. *Int. J. High Perform. Comput. Appl.*, 21(4):457–466, November 2007.
7. Erin Carson and Nicholas J. Higham. A new analysis of iterative refinement and its application to accurate solution of ill-conditioned sparse linear systems. *SIAM J. Sci. Comput.*, 39(6):A2834–A2856, 2017.
8. Erin Carson and Nicholas J. Higham. Accelerating the solution of linear systems by iterative refinement in three precisions. *SIAM J. Sci. Comput.*, 40(2):A817–A847, 2018.
9. Erin Carson, Nicholas J. Higham, and Srikara Pranesh. Three-precision GMRES-based iterative refinement for least squares problems. MIMS EPrint 2020.5, Manchester Institute for Mathematical Sciences, The University of Manchester, UK, February 2020.
10. Timothy A. Davis and Yifan Hu. The University of Florida Sparse Matrix Collection. *ACM Trans. Math. Softw.*, 38(1):1:1–1:25, 2011.
11. J. W. Demmel. *Applied Numerical Linear Algebra*. SIAM, 1997.
12. Jack J. Dongarra, Piotr Luszczek, and Yaohung M. Tsai. HPL-AI mixed-precision benchmark. <https://icl.bitbucket.io/hpl-ai/>.
13. I. S. Duff, A. M. Erisman, and J. K. Reid. *Direct Methods for Sparse Matrices*. Oxford University Press, second edition, 2017.

14. Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, USA, fourth edition, 2013.
15. 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.
16. Azzam Haidar, Ahmad Abdelfattah, Mawussi Zounon, Panruo Wu, Srikara Pranesh, Stanimire Tomov, and Jack Dongarra. The design of fast and energy-efficient linear solvers: On the potential of half-precision arithmetic and iterative refinement techniques. In Yong Shi, Haohuan Fu, Yingjie Tian, Valeria V. Krzhizhanovskaya, Michael Harold Lees, Jack Dongarra, and Peter M. A. Sloot, editors, *Computational Science – ICCS 2018*, pages 586–600. Springer International Publishing, Cham, 2018.
17. Azzam Haidar, Stanimire Tomov, Jack Dongarra, and Nicholas J. Higham. Harnessing GPU tensor cores for fast FP16 arithmetic to speed up mixed-precision iterative refinement solvers. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis, SC '18*, pages 47:1–47:11, Piscataway, NJ, USA, 2018. IEEE Press.
18. Azzam Haidar, Panruo Wu, Stanimire Tomov, and Jack Dongarra. Investigating Half Precision Arithmetic to Accelerate Dense Linear System Solvers. In *SCI6 ScalA17: 8th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems*, Denver, CO, 11/2017 2017. ACM, ACM.
19. Nicholas J. Higham. Iterative refinement enhances the stability of *QR* factorization methods for solving linear equations. *BIT*, 31:447–468, 1991.
20. Nicholas J. Higham. Iterative refinement for linear systems and LAPACK. *IMA J. Numer. Anal.*, 17(4):495–509, 1997.
21. Nicholas J. Higham. *Accuracy and Stability of Numerical Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, second edition, 2002.
22. Nicholas J. Higham. Error analysis for standard and GMRES-based iterative refinement in two and three-precisions. MIMS EPrint 2019.19, Manchester Institute for Mathematical Sciences, The University of Manchester, UK, November 2019.
23. Nicholas J. Higham and Srikara Pranesh. Exploiting lower precision arithmetic in solving symmetric positive definite linear systems and least squares problems. MIMS EPrint 2019.20, Manchester Institute for Mathematical Sciences, The University of Manchester, UK, November 2019.
24. Nicholas J. Higham, Srikara. Pranesh, and Mawussi. Zounon. Squeezing a matrix into half precision, with an application to solving linear systems. *SIAM Journal on Scientific Computing*, 41(4):A2536–A2551, 2019.
25. HPL. High performance Linpack benchmark. <https://www.top500.org/>.
26. *IEEE Standard for Floating-Point Arithmetic, IEEE Std 754-2008 (revision of IEEE Std 754-1985)*. IEEE Computer Society, New York, 2008.
27. Matrix algebra on GPU and multicore architectures (MAGMA), 2018. Available at <http://icl.cs.utk.edu/magma/>.
28. Intel. Math kernel library. <https://software.intel.com/en-us/en-us/intel-mkl/>.
29. Intel Corporation. BFLOAT16—hardware numerics definition, November 2018. White paper. Document number 338302-001US.
30. Heike Jagode, Asim YarKhan, Anthony Danalis, and Jack Dongarra. Power management and event verification in PAPI. In Andreas Knüpfer, Tobias Hilbrich, Christoph Niethammer, José Gracia, Wolfgang E. Nagel, and Michael M. Resch, editors, *Tools for High Performance Computing 2015*, pages 41–51. Cham, 2016. Springer International Publishing.
31. J. Langou, J. Langou, P. Luszczek, J. Kurzak, A. Buttari, and J. J. Dongarra. Exploiting the performance of 32 bit floating point arithmetic in obtaining 64 bit accuracy. In *Proceedings of the 2006 ACM/IEEE Conference on Supercomputing*, 2006.
32. Xiaoye S. Li and James W. Demmel. Making sparse Gaussian elimination scalable by static pivoting. In *Proceedings of the 1998 ACM/IEEE Conference on Supercomputing*, pages 1–17. IEEE Computer Society, Washington, DC, USA, 1998. CD ROM.
33. Software distribution of MAGMA version 2.5.0. <http://icl.cs.utk.edu/magma/software/>, January 2 2019. Quick reference <https://www.icl.utk.edu/files/print/2019/magma-sc19.pdf>.
34. C. B. Moler. Iterative refinement in floating point. *J. ACM*, 14(2):316–321, 1967.
35. NVIDIA. cuSolver library. <https://docs.nvidia.com/cuda/cusolver/>, November 28 2019.

36. NVIDIA Management Library (NVML), NVIDIA, 2018. <https://developer.nvidia.com/nvidia-management-library-nvml>.
37. US Department of Energy's Oak Ridge National Laboratory (ORNL). Summit supercomputer. <https://www.olcf.ornl.gov/summit/>.
38. Oliver Peckham. Summit achieves 445 petaflops on new "HPL-AI" benchmark. <https://www.hpcwire.com/2019/06/19/summit-achieves-445-petaflops-on-new-hpl-ai-benchmark/>, June 2019. Accessed July 22, 2019.
39. Youcef Saad. A flexible inner-outer preconditioned GMRES algorithm. *SIAM J. Sci. Comput.*, 14(2):461–469, 1993.
40. Youcef Saad and Martin H Schultz. Gmres: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. and Stat. Comput.*, 7(3):856–869, 1986.
41. Valeria Simoncini and Daniel B. Szyld. Flexible inner-outer Krylov subspace methods. *SIAM J. Numer. Anal.*, 40(6):2219–2239, 2002.
42. Robert D. Skeel. Iterative Refinement Implies Numerical Stability for Gaussian Elimination. *Mathematics of Computation*, 35(151):817–832, 1980.
43. G. W. Stewart. *Introduction to Matrix Computations*. Academic Press, 1973.
44. G. W. Stewart. The efficient generation of random orthogonal matrices with an application to condition estimators. *SIAM J. Numer. Anal.*, 17(3):403–409, 1980.
45. S. Tomov, R. Nath, H. Ltaief, and J. Dongarra. Dense linear algebra solvers for multicore with GPU accelerators. In *Proc. of the IEEE IPDPS'10*, pages 1–8, Atlanta, GA, April 19–23 2010. IEEE Computer Society. DOI: 10.1109/IPDPSW.2010.5470941.
46. Stanimire Tomov, Jack Dongarra, and Marc Baboulin. Towards dense linear algebra for hybrid GPU accelerated manycore systems. *Parallel Computing*, 36(5):232 – 240, 2010.
47. J. H. Wilkinson. *Rounding Errors in Algebraic Processes*. Prentice-Hall, 1963.